TTI VANGUARD
ADVANCED TECHNOLOGY RESEARCH MEMBERSHIP

BLOCKCHAIN AND OTHER NETWORKS HIGHLIGHTS

September 12–13, 2019 • Washington, D.C.

prepared by Nancy Kleinrock nk33@cornell.edu

Blockchain-Based Payment Systems—Dr. Bradley Chase, Ripple

- Just as the underlying purpose of the Internet is to facilitate the seamless movement of data, Ripple's vision is to facilitate the seamless movement of value—notably in the form of payments—regardless of national boundaries.
  - Contrast this with the prevailing reality of international payments: a slow, expensive, unreliable, and cumbersome process.
  - Chase describes this process and how the XRP blockchain protocol—and the ecosystem of financial services that has grown up around it—is opening the door to freer exchange that stands to ignite business opportunities and open up markets.
  - "I can send an email around the world; why can't I send $100 around the world?" poses Chase.
- To understand the need, Chase depicts an international-payment scenario:
  - Suppose Bob is working up a presentation and enlists the services of graphic designer Alice.
    - While discussing his needs, Bob wants to purchase Alice a cup of coffee as a token of appreciation.
    - If their consultation is taking place face to face in a coffee shop, Bob walks to the counter and graciously pays cash for her beverage.
    - But what if the consultation takes place over Skype, with Bob in South Africa and Alice (and her coffee) in the Philippines?
    - A money transfer service would add a sizeable surcharge, involve the uncertainty of fluctuating monetary exchange rates, and require several days to settle via S.W.I.F.T., meanwhile tying up funds; and that is only if all goes well—6% of transactions fail.
      - The friction arises from the banking system, which is laden with legacy systems, regional rules, bilateral relationships between institutions, and more.
      - S.W.I.F.T. is moving toward faster settlement, while working with banks on interfacing with their legacy systems and/or encouraging them to modernize; banks in favor of this change tend to be similarly encouraging of Ripple's efforts.
        - Nevertheless, S.W.I.F.T. remains reliant on its closed, one-way messaging scheme.
    - Alas, it looks as though Alice will have to buy her own coffee.
    - "It is not just one bottleneck that causes the challenge, but is death by 1000 cuts," says Chase.
- Ripple addressed this conundrum by applying blockchain to the portion of the transaction involving international exchange by providing fast, cheap, decentralized settlement of the (digital) asset.
  - Chase defines blockchain as a specific type of data structure: a sequence of states of a distributed database linked by cryptographic hashes.
    - For a permissionless public blockchain like Ripple, all state, all transactions, and all transaction rules are public, and all participants can enforce rules and verify the state of the database, and thereby verify its integrity.
    - To be effective in the relevant context, Ripple defines the distributed state that is the blockchain to represent a distributed ledger associated with the distributed asset that is the XRP cryptocurrency.
      - "The digital asset is intrinsic to the ledger," says Chase. "It is only grounded in that ledger itself."
    - The transfer of funds, such as from Bob to Alice for coffee, is effected by Bob using his private key to sign the transaction of transferring a specified portion of funds allocated to his public key to Alice's public key.
      - Transaction rules enforce (a) no double spending and (b) a specified supply schedule, as dictated by the order of writing new transactions as subsets of blocks written to the blockchain.
    - It is the decentralized nature of Ripple's blockchain that will let anyone transact with anyone else.

      - As such, participation is open and permissionless, and the system is resistant to censorship, bears no administrative functions, and is fault tolerant by design, even in the assumed presence of malicious participants.
        - “Anyone who is willing to participate and abide by the rules, can have their transactions submitted.”
  - What has been described thus far could also pertain to the process of the Bitcoin or Ethereum networks; an essential difference among these and Ripple is the consensus protocol.
    - Bitcoin and Ethereum both use lottery-based leader selection to determine the next block to add to the chain and, thereby, the order of transactions (Bitcoin uses proof of work, while Ethereum is migrating to proof of stake).
      - Any lottery-based system benefits from minimal coordination, but suffers from a trend toward centralization (e.g., for miners advantaged by great computational power, in the proof-of-work scenario) and lack of determinism (e.g., in the event of nearly simultaneous puzzle completion (proof-of-work) or near-equal stake).
    - XRP is instead a Byzantine consensus protocol rooted in trust-based cooperation.
      - This class of protocols relinquishes a degree of coordination in favor of determinism and utility, as measured by rapid addition of each new block to the blockchain at minimal computational cost.
        - By design, a new block is appended to the Bitcoin blockchain roughly every ten minutes, while Ripple’s expands every three seconds or so.
      - Cooperation in the XRP ledger consensus protocol is expressed through the mutual selection of a set of peers who propose candidate transaction sets, which are collaboratively refined until convergence is reached and a set is adopted as the next block to add to the ledger.
        - Robustness is ensured since any malfeasance that might play out would be readily detectable by other participants, resulting in the booting of that individual from the network.
        - Although cooperation is an inherent facet of the XRP protocol, currently there are upwards of 150 known validators of the XRP ledger, with Ripple itself operating only a handful itself.
    - Summarizing, compared to Bitcoin and Ethereum, XRP is as much as three orders-of-magnitude cheaper and faster.
      - Other benefits are that accounts on the XRP ledger are primary objects, enabling accounts to own a set of private keys or issue assets on the ledger, engage in payment channels with high-volume–low-value transactions that transact offline with periodic on-chain settlement, or establish a cryptographically secured escrow with funds that only transfer after the consummation of a smart contract.
- The Ripple model removes friction in cross-border transactions by eliminating the need for banks in different jurisdictions to engage bilaterally; instead, each interfaces solely with the Ripple exchange: buy XRP at the source exchange (e.g., South Africa), transfer units in a few seconds to the destination exchange (e.g., the Philippines), and sell them there for pesos.
  - “The power comes from separating foreign exchange from crossing the border,” says Chase.
  - Not only does this avoid the need for foreign exchange, but it also enables liquidity on demand, with no need for prefunding, provided many participants exist in each jurisdiction.
  - “Network effects grow the liquidity for you,” he says, while also emphasizing 24/7/365 availability of the exchange.
- Despite XRP’s beneficial attributes, they will only truly shine when Ripple’s model breaks out of its own sandbox and is capable of interoperating with the broader ecosystem related to the transfer of value; that is, “A single blockchain cannot solve the entire payment experience,” says Chase. “Interoperability is essential.”
  - Toward this end, Ripple proposes Interledger—a protocol for connecting ledgers.
    - Interledger serves the same role in the value transfer architecture that IP serves in the Internet architecture; that is, it sits at the neck of the hourglass.
    - Just as IP integrates networks and sends data because of its open standard, use case agnosticism, and reliance on packetization of data, the Interledger protocol (ILP) establishes interoperability and sends value because of precisely the same feature set (replacing packetization of data with packetization of value).
    - Specifically, any given payment is chopped into micropayments, with average value in the millionths of cents.
      - If some micropayments fail, no great loss is suffered; similarly with packetized data.

- With this architecture, the need for quotes, options, and trust disappear, and horizontal scalability emerges naturally.
- As Interledger connects the growing array of digital payment systems, the protocols associated with each such blockchain-based value instrument become more robust, leading to a vibrant and reliable Internet of Value.

- Although Chase's focus has been on the utility of the XRP protocol and ledger to ease the process of foreign transactions, it is also true that XRP—like other cryptocurrencies—serves as a speculative asset and therefore its price will likely continue to experience volatility until its functional use dominates.

Deconstructing Blockchain—Mr. Roger Meike and Mr. Glenn Scott, Intuit

- As is the case with many firms, Intuit has been exploring how blockchain technology might best enhance its business—specifically, providing financial and tax preparation software to businesses and the public.
  - As director and key engineer, respectively, of Intuit's innovation team, Meike and Scott decided to go beyond reviewing blockchain solutions developed or adopted by others, and instead embarked on a journey to deconstruct the blockchain into components, explore the features of each, and only then determine the best path forward.
  - In the end, they deemed that a blockchain per se was not optimal, but instead could be used to inspire the development of the internal tool Blackflower, which weds the immutable ledgering of blockchain with an object-orientated mindset.
  - Meike, as the team's mouthpiece, presents an Alice-and-Bob scenario to introduce the ins-and-outs of blockchain—specifically, the Bitcoin model—and then uses this as a jumping off point for a brief discussion of Blackflower.
- A blockchain story of Alice and Bob (and Carol and Ted):
  - Bob is in a bad way: He has an inconsiderate roommate, Carol, who is driving him to distraction—so much so, that—with eight months yet on their one-year lease—Bob takes to the dark web, where he unearths professional hitman Alice (aka Cooper) and contracts with her to do what she does best.
    - Not surprisingly, she is unwilling to accept traceable payment (e.g., credit card or bank transfer), nor is she willing to meet Bob face to face to be paid in cash; instead, bitcoin is her currency of choice.
  - Bob initiates the transaction by inputting his "address" (i.e., his public key) and proceeds by specifying Cooper's address, the amount in bitcoin to transfer to that address, along with a tip for the bitcoin miner; Bob signs this wholly pseudonymous transaction with his private key.
    - Note the lack of mention of the balance of bitcoin in Bob's account; only the amount of each transaction enters the blockchain, and the balance must be computed by summing all inputs associated with his address and subtracting from that the total of all outputs.
    - To verify one's balance, an individual must either rely on a service or targeted software to search the full blockchain for all transactions associated with their respective cryptographic address.
  - After Bob's transaction is broadcast to the Bitcoin network, miner Ted, who operates a large server farm in China, combines Bob's transaction with a large number of other transactions to create a block and competes with other miners to have his (Ted's) block accepted to the blockchain.
    - Bitcoin's proof-of-work competition entails finding a hash of the block that contains a requisite number of leading zeros.
      - This process begins with the hash of the most recently added block and adds to it a nonce and the set of newly accumulated transactions, including one transaction unique to that miner's block, namely self-payment that will proceed only if that block wins the competition and is added to the blockchain.
        - The "work" amounts to successively incrementing the nonce until either the hash achieves enough leading zeros or until a different miner's hash does so and becomes accepted as the next addition to the blockchain.
          - Note that a small change in the input to a good hash function will generate a wildly different hash; therefore, the presence of each miner's self-payment transaction guarantees that each will be working on a distinct computational problem.
        - The number of requisite leading zeros is recomputed once every two weeks, depending on the volume of bitcoin mining, to ensure that a new block is added to the Bitcoin blockchain roughly every ten minutes.

            - Once a miner declares victory, they broadcast the winning block to the Bitcoin network, where other miners validate the block's transactions by confirming each signature and computationally verifying that no double spending has taken place.
            - As each node in the Bitcoin network accepts the new block, that node moves on, restarting the work of competing to have their block win the next round.
        - While hash chains have been well known since the 1980s, the true innovation of Bitcoin, believes Meike, is its incentive structure.
            - Self-payment comes in the form of newly minted bitcoins—hence the mining metaphor.
            - The amount of bitcoin created with the adoption of each new block was structured from the outset to decrease over time to avoid runaway inflation in the cryptocurrency.
            - Initially supporting—and eventually supplanting—payment in new bitcoin are tips paid by those, like Bob, who submit transactions; the larger the tip, the more likely miners will accept a given transaction into its candidate block.
            - Combined, this incentive structure has made the Bitcoin network large and vibrant.
            - On the downside, the intrinsic electricity-intensive nature of bitcoin mining is proving to have deleterious effects in its own right; there exist other consensus algorithms without this downside.
    - Returning to our cast of characters, once Ted (or another miner) succeeds in having his block—including Bob's payment to Alice—accepted to the blockchain, the payment in bitcoin is now Alice's, with no take-back opportunity by Bob.
        - As Alice begins to make her nefarious move on Carol, she instead finds him appealing—not obnoxious—and doesn't off Carol but rather begins to date him.
        - The happy ending to the story is that Carol moves out of Bob's place (he breathes a sigh of relief at Carol's departure) and into Alice's upgraded home, thanks to her income from Bob.
- Some general comments about blockchain technology:
    - The blockchain is immutable; once a block has been added, neither it nor its contents can be altered.
        - Generally, this is considered a positive feature, but some potential blockchain applications require reversible transactions; consider, notably, GDPR's right to be forgotten.
    - Transactions are pseudonymous; public keys are exposed, but not the names associated with them.
    - The presence of each transaction in the ledger is proof of its existence: "The transaction *is* the ledger entry *is* the receipt," says Meike. "It is like swearing to the world that the transaction has taken place.'
    - The blockchain ledger serves as an independent transaction record and thereby makes possible an additional crosscheck on payments.
        - This "triple-entry accounting," recorded in a public, cross-cutting ledger, could usher in transactional trust across the entire economy.
    - The blockchain ledger tracks ownership, by design; currency is hardly the only thing that is owned, suggesting an unending array of applications.
    - The Bitcoin network is an experiment in a decentralized monetary system—notably, one that is not tied to any government.
        - To disassociate the technology from Libertarian ideals, many have taken to referring to blockchain instead as distributed ledger technology.
    - Of the nearly 2000 distinct cryptocurrency specifications currently deployed, many do not use the Bitcoin incentive structure verbatim; success has been spotty at best.
    - While the Bitcoin blockchain is open (in the sense that all transactions are visible, albeit using pseudonyms rather than names to identify them), many blockchain implementations are instead permissioned, such that only authorized participants are permitted access; loss of visibility adds privacy, but can lead to loss of trust.
    - The Bitcoin blockchain relies on a proof-of-work consensus algorithm (mining), but other blockchains rely on proof of stake or a trust-based Byzantine consensus protocol.
    - Many blockchains exist to support transactions in bespoke cryptocurrencies, but Ethereum distinguishes itself with "smart contracts" that associate transactions on the blockchain with the ability to run pieces of code and thereby enable a broad range of applications.
    - Despite careful planning and design, it is impossible to envision every anomalous situation and preestablish a resolution for it; as such, governance is an ongoing issue facing blockchain networks.
- Intuit's Blackflower is not a blockchain implementation per se, but rather looks to aspects of the blockchain ethos for inspiration.
    - Reasons against blockchain at Intuit:

  - Over the past 35 years, the firm has built a brand its customers trust; why, then, adopt blockchain wholesale and relinquish that trust to a technological solution?
  - Privacy regulations require the potential to remove personal data at an individual's request; their immutability makes blockchains ill-suited to such applications.
- Yet ledgers are convenient receptacles for information that changes over time, whether conventional transactions (where ownership of assets ebb and flow) or attributes of, say, an individual customer.
- As such, Meike's team has chosen to "go small" and record the history of each customer in an individualized hash chain that it dubs a "journaled object."
  - Each object is an independent, immutable ledger that represents not only the current state of a person, business, or property, but through journaling it also reflects that object's full history.
  - Drawing two or more objects together into a combined journaled object represents a relationship among the distinct entities.
    - A relationship might express ownership (person and asset), employment (person and employer), or trade (transfer of asset from one person to another).
      - For instance, Alice and Bob—each with their own representation as a journaled object within Intuit's system—unite in a shared transaction in which Bob makes payment to Alice in exchange for goods or services; their independent ledgers update with the new information resulting from their transaction.
        - If Alice decides to retract all of her data from Intuit, with it will go this joint journaled object, but Bob's individual ledger would retain evidence of payment (the transaction does not disappear, only Alice's role as a partner in it).
          - On an object-by-object basis, one or more aspects of a relationship-based journaled object (e.g., a transaction) could be relegated to a public blockchain for verification and safekeeping.
        - Each journaled object is immutable, due to being rooted in a hash chain, but an entire chain can be expunged from the system.
  - Meike recognizes that this journaled-object construction is a type of an append-only database, but one with an interesting degree of granularity.
  - Building on Intuit's reputation, it serves as the single source of truth in this information ecosystem.
    - With the lightweight structure of journaled objects, transactions can be instantaneous.
    - A tax-preparation example: Carol has moved, and enters his new address (i.e., Alice's) into TurboTax. He now lives in New Jersey, across the Hudson River from Bob's place.
      - Carol's journaled object, recording this change, triggers the tax-prep software to pose questions about needing to file in more than one state (yes) and whether his marital status has changed (not yet—and perhaps not ever, for Alice is just discovering some of what Bob had come to know, namely that Carol leaves piles of sweaty workout clothes in a damp heap, drinks milk directly from the carton, and never washes the dishes).

Enterprise Blockchain: Case Studies, War Stories, and New Opportunities—Mr. Brian Behlendorf, Linux Foundation

- When a new technology hits the scene that features not only multifarious potential uses but also multifarious implementation strategies, the potential arises for a platform war to erupt.
  - This serves no one—well, no one aside from the ultimate winner (think VHS vs Betamax, with both ultimately losing out to first DVDs and then streaming).
  - Fortunately, when it comes to software, there is a long history of collaboration among competitors to improve the offerings; this, of course, is the open-source software movement, now 20-years-old and always improving.
  - The Linux Foundation has long been the beating heart of open source, first with its operating system, and more recently with its ever-expanding stable of projects, consortia, and educational initiatives.
    - With software underpinning every facet of society, it is no surprise that the Linux Foundation serves as an umbrella under which the most powerful companies in the world unite to improve the software infrastructure for domains as distinct as security, networking, cloud, automotive, the Web, blockchain, and beyond.
      - "We are kind of like the nerd World Economic Forum," says Behlendorf. "We serve as this convener for all sorts of technology projects, well beyond blockchain technology and well beyond the Linux operating system. We are arguably the largest shared technology investment in history."

  - Regarding blockchain, the project name is Hyperledger; it was founded in late 2015 with an aim to support the collaborative advancement of blockchain-based distributed ledgers with open-source tools.
  - Behlendorf describes the project and the tools that it has spawned—most notably Hyperledger Fabric—as well as describing a set of case studies rooted in this technology.
- In 2015, as cryptocurrencies were blooming like goldenrod in early autumn, it became clear that the time was ripe to establish some clarity in the blockchain space.
  - With the Bitcoin network well established as the exemplar of blockchain, it was clear that the platform's electricity-intensivity would contribute to the downfall of the planet if similar proofs of work became the standard across all blockchain networks.
    - "Burning all the CPU power in the world to win a lottery to decide who gets to put the next link in a data structure was the most anti-environmentalist thing I could do," says Behlendorf. "It felt just wrong, and the way it had been used by speculators driving pump-and-dump schemes made a lot of enterprises very nervous about it."
  - Yet the practical promise of blockchain technology as a distributed ledger to converge on a common source of truth and to certify transactions, with no need for a trusted third party who would control the platform and extract rents from all its users, struck Behlendorf as a noble endeavor to support.
  - "The formative concept behind Hyperledger was to build communities of business participants in these blockchain networks, using blockchain technology to transform these cross-business workflows," he says. The will was there, but there was not then a depth of thought of what made "a proper blockchain network."
    - Industries each have their own set of member types and their own set of interorganizational activities, but the underpinning structure of a blockchain implementation can be common to all.
      - Examples: financial services (e.g., bank wires, equity trading, mortgage underwriting, P2P lending); supply chain (e.g., provenance tracking, trade, finance, customs, IoT asset tracking, title tracking); healthcare (e.g., provider directories, provider certification, permissioned health record sharing, insurance claims, pharma supply chain).
    - Just as the underpinning technologies of the Internet enabled a network of networks, Hyperledger's suite of ledgers and associated support structure enables players within a given industry to work together efficiently, trusting the ledger to maintain the ground truth and accurate history.
    - "I am convinced that we won't see the majority of transactions on public ledgers," says Behlendorf. "They will be fanned out into these permissioned ledgers, and some of them will be tiny, with ten participants, and some of them will be very large, with thousands of participants, even within that permissioned model."
      - When considering permissioned blockchains, permission to write is a more stringent constraint than permission to read.
- Hyperledger currently hosts a "greenhouse" of distributed ledgers, libraries, tools, and domain-specific solutions, but all share the same software license and the same software development methodology—one that forces industry competitors, now acting as collaboration partners, to work together toward a common goal while also determining what components to separate out for independent proprietary development.
  - Hyperledger Fabric, originally contributed by IBM, is the most broadly deployed enterprise platform; 32 of Forbes' "blockchain 50" firms cited Fabric as an essential consortia ledger when queried about the technologies they rely on.
    - Fabric is but one of Hyperledger's distributed ledgers; it is joined by Besu, Burrow, Indy, Iroha, and Sawtooth, each with a targeted focus.
  - The governing board of the Hyperledger project has members from some 20 firms, with many dozens of global organizations contributing the hard work necessary to the project's success.
    - "You have to have everyone in the room to make this have any credibility," says Behlendorf.
- Case studies:
  - IBM Food Trust—Partnering notably with Walmart, Unilever, Nestlé, Carrefour, and Dole, IBM has applied blockchain technology to the food supply chain to enable visibility by all permissioned parties on a real-time basis (two-second tracking) into the provenance, handling, regulatory compliance, and other pertinent facets of how food is managed throughout the supply chain from farm to grocery store.

    - The first application was to leafy greens in an era when a single farm's *E. coli*-tainted spinach could be distributed to over two-dozen states, leading to a nationwide recommendation to avoid eating fresh spinach altogether.
    - The power of the Food Trust network is not only the fast-tracking of food items, but also avoiding concentrating all the data—and all the power and need for trust—with a single stakeholder (e.g., Walmart).
    - Carrefour reported that it has recognized some $2M in increased revenue due to consumer preference for traceable food.
- Chinese banks and letters of credit—This endeavor has provided the Chinese banking industry with a jointly developed blockchain-based platform for transmitting letters of credit that is open, standardized, and compliant with regulations.
    - Letters-of-credit transactions amount to roughly 1B RMD ($150M) daily, featuring shortened delivery time and improved efficiency and security.
    - This platform has been enthusiastically embraced by 20 banks in an industry that had grown increasingly frustrated by its paper- and fax-based predecessor.
    - A key contributor to success has been the Chinese-language, business-specific user interface.
    - Consortia ledgers are especially valuable in parts of the world with weak, corruptible, or inefficient governmental structures, where it is incumbent on industry partners to assure best practices and workflows.
- Eshare—Also native to and serving China is the JD blockchain open platform Eshare, which uses Hyperledger Fabric to improve the process of resolving disputes regarding digital rights.
    - Customers can gather digital evidence—notably electronic invoices—from the blockchain invoice-and-contract platform and fast-track appeals directly with the Guangzhou Internet Court, which in turn relies on the blockchain to verify integrity.
    - "Adjudication processes that once took months to years now only take days to weeks," says Behlendorf.
    - While assertions on the blockchain could be fraudulent, the blockchain itself will natively expose some aspects of fraud, such as double spending, leading perpetrators to being kicked off the network.
        - "The network can enforce the antifraud mechanism and keep it from happening in the first place," he says.
- Change Healthcare Intelligent Healthcare Network—This Hyperledger Fabric-based network supports the fast and efficient routing and settling of insurance claims.
    - How fast? 50M transactions daily; 550 each second—sufficient to handle all U.S. healthcare claims.
    - Previously, Change had operated as a centralized nexus of healthcare claims, but found itself being shunned by certain members of the ecosystem that did not want to trust one firm in that role.
    - Now, any authorized party among the 20 participating partner organizations, whether a hospital administrator or government payer, can view the full history and real-time status of all transactions relevant to that entity.
- OrgBook—OrgBook is the easy-to-join, secure, interoperable, trusted digital network of verifiable data based on Hyperledger Indy that the government of British Columbia built to enable business and governmental entities to assess evidence that a potential business partner is legally incorporated.
    - An easy enrollment process yields each participant a Self-Sovereign Identity that is stored on the public, globally accessible blockchain network.
    - This fraud-reduction service reduces bottlenecks in information gathering while also improving privacy.
    - Example use case: A properly incorporated business exposes its digital identity variously to the restaurant board, the alcohol authority, and other relevant agencies to efficiently receive the legal go-ahead to launch operations.
        - "You, as the business owner, are the pivot point for all of these relationships," says Behlendorf. "You have these documents that cannot be taken away from you; they might have expiration dates, but you can take them and present them to other entities to tell them that you have the right to serve alcohol."
    - A generalization of OrgBook might serve individuals who register on the blockchain and use it as a generalized sign-in identity, revealing personal data (e.g., driver's license, passport, birth

        certificate, components of healthcare record) on a need-only basis, with permissions specified by that person.
    - Kiva's Digital ID and privacy-first credit bureau for Sierra Leone—To bolster Kiva's ability to meet the microlending needs for Sierra Leone, it worked with the government to develop and launch a Hyperledger Indy-based digital-credentials network.
        - Moreover, the nation's central bank will use Hyperledger Fabric as the technological underpinning of a shared, decentralized credit-reporting bureau.
        - With these pieces in place, Kiva will be better situated to connect with, vet, and lend funds at reasonable/humane rates for the small projects that can change the lives of those on the receiving end of such loans.
        - "They have now enrolled 8.5M IDs and are starting to provision these, biometrically linked," says Behlendorf. "All of these data are kept off the blockchain. What is on the blockchain are the keys, the signatures, and hashes that are integral to enabling somebody to prove—even in a low-resource environment—who they are and prove their history of financial transactions, and use that as the basis for being able to obtain credit at nonusurious rates."
- With the title of his presentation promising war stories, Behlendorf alludes to—but does not flesh out—some challenges faced by the major cloud providers when implementing Hyperledger as a hosted service.
    - Governance of blockchain networks stands out as a matter requiring further attention.
        - "Governance models matter here," he says, "and that's the technical governance, but also the human governance and the corporate governance between these different parties. We will find templates and repeatability there, I am quite sure."

Privacy-Preserving Technologies Meet Machine Learning—Dr. Jeanette Wing, Columbia University

- Blockchain is hardly the only privacy-preserving technology; to put blockchain into perspective, when Wing turns her attention to privacy issues related to machine learning, the toolbox she uses does not include the titular technology of this conference.
    - Instead, she describes a pair of cryptographic strategies—secure multiparty computation and homomorphic encryption—along with the statistical approach of differential privacy.
    - Most powerful is a combination of these approaches, including the beneficial introduction of noise that underpins differential privacy to eliminate the potential of doing a successful diff against the results of a secure multiparty computation to reveal the participating parties.
    - Wing describes this tactic as well as others at the intersection of machine learning and privacy, emphasizing that there is no one-size-fits-all solution to preserving privacy, but rather the need to find solutions that best fit the problem at hand.
        - "There is value we can all gain when we can pool our datasets, and we see this in spades with machine learning," says Wing. "How can we share data while preserving privacy? There are technical solutions to these problems—not in general, but for point problems there can very well be a point solution."
    - She appends her remarks by introducing the joint partnership between Columbia University and IBM: the Center on Blockchain and Data Transparency.
- Cryptographic technologies:
    - Secure multiparty computation:
        - In 1982, Andrew Yao proposed a technique to preserve the privacy among participants in a joint computation.
        - The model system—the so-called millionaires' problem—entailed a pair of rich folks (Alice and Bob, of course) wanting to know who was wealthier, but without divulging how much dough each was rolling in and without enlisting a trusted third party to do the comparison and report back.
            - The strategy involves garbled circuits, where Alice—the garbler—garbles (e.g., encrypts) the Boolean circuit that describes the comparison function, sending this as well as her directly encrypted input of her wealth to Bob; Bob decrypts the circuit, and the two compare the encrypted outputs to learn the result of the comparison, although not the encrypted values.
            - The crux of the process is the encoding of the comparison function as a garbled circuit.
        - This technique has since been generalized from two to multiple parties and has been applied to, for instance, the Boston wage equity study (a rather direct extrapolation of the millionaires' problem), the Danish sugar beets auction, and research associated with genomics, social media behavior, ad targeting, and network monitoring.

- Microsoft and Google are known to use secure multiparty computation to perform the cryptographic process of private set intersection to compare the overlap in their customer base.
  - Alice's firm has the set of customers $X = \{x_i\}$, while Bob's firm has the set of customers $Y = \{y_i\}$; what is the overlap set $O = X \cap Y$? The answer reveals to both their joint customers without spilling the beans on customers of one but not the other.
    - Pointing to relevant papers without going into procedural details, Wing notes that Microsoft declares scalability up to 1B customers, while Google touts their ability to bring information about money into the comparison, as well.
- Secure multiparty computation becomes practical in the machine-learning space when information from more than one source serves as input to a joint machine-learning endeavor, yet that input must be kept private from all but its source.
  - Example: Patient databases from multiple hospitals pool to form the necessarily large training set from which a third-party generates a machine-learning model that can then be used to evaluate whether a given patient's data suggests the incidence of a particular disease.
    - Garbling the data input to the pooled repository and garbling patient-specific output from the model satisfies patient privacy regulations.
    - "You have this garbled circuit that takes this garbled data," says Wing. "Then you can query this super-duper machine-learning model and give back the result, such that the individual clinics can decipher the individual results and determine whether a particular patient might have cancer, for instance."
- Homomorphic encryption:
  - A breakthrough in fully homomorphic encryption—the ability to compute over encrypted data—came with Craig Gentry's 2009 dissertation on lattice-based cryptography, with subsequent improvements that reduce the computational load, although feasible implementation remains out of reach for many applications.
  - It has, however, proven useful in the context of cloud computing, where a user doesn't trust the cloud provider and therefore encrypts their data before uploading it; in principle, arbitrarily sophisticated computation can then be performed over that encrypted data, returning the result in encrypted form, which the user decrypts with the key used for initial encryption.
    - Homomorphism, in this context, states that the result of encrypting *a*, encrypting *b* and summing the result ($E(a) \oplus E(b)$) is equivalent to encrypting the sum of *a* and *b* ($E(a+b)$).
    - Gentry proved that this relationship holds, theoretically, for any Boolean or arithmetic circuit, although the computational difficulty compounds as the complexity of the circuit intensifies.
  - In practical terms, the use of homomorphic encryption in machine-learning applications boils down to restricting the nonlinear functions of deep neural networks to simple functions.
    - The simplest possible nonlinear polynomial is the quadratic x^2, which Microsoft researchers employed in a five-layer neural net evaluating the MNIST handwriting dataset using the Simple Encrypted Arithmetic Library.
      - "It worked. You get the speed of this very fast SEAL library using x^2 instead of a sigmoid, and you get entire computations in the homomorphic-encrypted way, with prediction results on the order of seconds, and it is highly parallelizable," says Wing.
  - Example: Successful and "fast enough" data protection using homomorphic encryption when using a machine-learning model to analyze genomic data in the 2015 NIH Genome Analysis Competition.
- Statistical technologies:
  - Differential privacy:
    - "Anonymization doesn't work, so don't rely on preserving the privacy of individuals based on anonymization," says Wing, citing several infamous cases in which anonymized data, combined across data sources, permitted the re-identification of individuals.
    - Differential privacy comes to the rescue by adding properly constructed noise (e.g., with a Laplacian distribution, but not Gaussian) to data such that output derived from the dataset cannot distinguish whether a particular individual's data is or is not present in the dataset.
      - That is, deleting a record from the database generates the same result to a statistical query over the database as would be the case were that record retained.
    - Specifically, differential privacy is a stability constraint on computations running on datasets that requires that no single data point in an input dataset has significant influence on the output.
      - "A small perturbation of the input is, within some parameter, undetectable," says Wing.

      - It is the addition of the noise that confers anonymity, rather than stripping out overt identity information, such as name, birthdate, and so forth.
    - Example uses, which demonstrate the use of this technique by technology’s titans:
      - Google has used differential privacy in its SafeBrowsing malware detection tool to gather statistics on harmful websites without tying individuals to their browsing histories.
      - Apple uses differential privacy for iPhone data.
      - Facebook has expressed interest in differential privacy for advertiser analytics.
      - Microsoft uses it internally for data analytics.
      - Uber uses it for traffic analysis.
      - In 2020, the U.S. Census Bureau will use differential privacy to protect post-collection datasets from being used to re-identify anonymized data.
        - “By law, the Census Bureau has to guarantee the people who are being counted that their identities will not be revealed,” says Wing. “Differential privacy is not the salvation, but it is all that there is; it is the best we can do.”
        - Nevertheless, in the presence of noise, social scientists will be faced with the “interesting research challenge” of restructuring their census-based analyses.
- Combining cryptographic and statistical technologies:
  - Returning to the machine-learning-over-healthcare-data example above, incorporation of differential privacy into the garbled results delivered to each hospital following secure multiparty computation over the model learned from the pooled dataset renders ineffective any attempt to do a diff to deanonymize data.
    - “Before sending back that result, add a little noise,” says Wing. “Then you can’t do a meaningful diff.”
  - Alleviation of a possible concern:
    - It is well known that image-oriented deep neural nets often fail catastrophically in the face of small perturbations; e.g., fail to properly recognize a traffic sign if a snippet of duct tape obscures a bit of it.
    - Adversarial machine learning is the current strategy of choice to improve outcomes.
    - Wing does not find it useful to denigrate techniques based on outlier bugs that can defeat them; instead, she sees value in proving that adding a noise layer to the deep neural network—à la differential privacy—guarantees that the classifier will be robust to a specifiable degree to input perturbations.
      - “It works really well if you insert the noise into one of the early layers,” says Wing.
- Joint partnership between Columbia University and IBM—The Center on Blockchain and Data Transparency:
  - Built on Hyperledger technology, this center will encompass an innovation accelerator as well as research and educational initiatives with focus areas including identity management, valuation, ethics, privacy, quality, governance, and trustworthiness.
    - “The goal is to cut through the hype to discover what is blockchain, what does blockchain actually provide you, and what guarantees can you actually provide customers,” says Wing.

Lessons from the Blockchain Impact Ledger—Ms. Dahna Goldstein, New America

- There are a dizzying variety of ways that existential challenges impact people around the globe.
  - But so too are people and organizations striving to make a positive social impact.
  - While those living on the edge might not have blockchain front of mind, those on the assistive end increasingly do.
  - With its Blockcahin Impact Ledger, the Blockchain Trust Accelerator—a platform within New America—hosts an “ironically centralized” go-to repository of projects that rely on blockchain technologies to support the work of organizations aligned with the 17 social development goals (SDGs) defined by the United Nations.
    - “We also bring together governments, technologists, civil society organizations, and philanthropists who are interested in exploring the potential of blockchain and other technologies to improve their interactions with their constituents, to increase accountability, to do all of the things that blockchain has the unique potential to do,” says Goldstein.
      - Some potential impacts of blockchain that Goldstein envisions: improve governance, reduce corruption, improve aid delivery, secure voting, improve sustainability, legitimize property rights, change the future of work, and create opportunity.

            - No small fraction of her work has been cutting through the negative hype associated with blockchain, largely stimulated by the high volatility and nefarious Dark Web uses of cryptocurrencies.
    - Goldstein first describes the Blockchain Impact Ledger itself and then presents several projects included on it.
- The Blockchain Impact Ledger is a resource for entities looking to leverage the power of distributed-ledger technology for good.
    - To be included in the Ledger, a project must (a) feature some degree of scale, (b) confer social impact, and (c) demonstrate overt upside from including blockchain to achieve stated social impact goals.
        - Each of these three criteria are, on the surface, difficult to measure and assess, making it incumbent on Goldstein's team to develop a methodology for vetting projects.
        - "Many of the projects are still in the early stages."
- Case studies:
    - Remittances—Ant Financial, BitPesa
        - Given blockchain's origins in the cryptocurrency sphere, several of the projects contained in the Ledger involve the transfer of funds, notably for remittances; Ant Financial is one of these.
            - Conventionally, the middlemen inserted into the transfer of remittances strip out a significant fraction of funds for the service they provide; a peer-to-peer donation-facilitation tool creates positive social impact by delivering the full remittance amount to the recipient, helping to lift up that individual or family from poverty.
            - As specified in the Blockchain Impact Ledger, this private, proprietary blockchain partners with GCash and Standard Chartered to "improve the speed, efficiency, safety, and transparency of cross-border remittance"; the project is aligned with the SDG of reducing inequalities.
        - Another such project is Nairobi-based BitPesa, which has been facilitating the transfer of money (in the form of bitcoin) within Africa in a fast, easy, and cost-effective manner since 2013 and is currently serving as many as 100K people.
    - Reducing food insecurity by providing funds for refugees to purchase food—UN World Food Programme's Building Blocks
        - Instead of distributing food directly to refugees who have enrolled using biometric identity validation (since some lack papers), the World Food Programme provides funds directly to recipients through an Ethereum-based private blockchain, reducing transaction fees compared to conventional options; these funds can then be used to purchase food through a network of vendors.
            - Although the UN food bank is the sole administrator of the blockchain, it is populated with all the vendors that sell into the food distribution network as well as each consumer-facing grocer.
            - "They are now writing all of the transactions onto a blockchain and have reduced transaction costs by 98%, saving $40K/month," says Goldstein.
                - The savings derive from the World Food Programme now working with one bank, rather than the myriad banks with which individual food vendors previously partnered.
        - This project is meeting the SDGs of reducing hunger and engaging in partnerships to achieve goals as it serves more than 100K individuals.
    - Supply chain management—Moyee Coffee, Everledger Diamond Provenance Platform
        - Blockchain is increasingly becoming a staple technology for supply chain management.
        - In the case of the Ethiopian coffee company Moyee, not only does the Hyperledger-based blockchain document where beans are harvested, roasted, and packaged—with this information visible to consumers when scanning the QR code on the package—but also enables the consumer to tip the farmer with a digital token, closing the loop between consumer and producer.
            - The objective of the project is to create transparency and integrity in the coffee supply chain, and the relevant SDGs are listed as zero hunger, and responsible consumption and production.
        - Everledger uses a private blockchain (IBM Blockchain Platform and Hyperledger Fabric) to track the provenance of diamonds by evaluating compliance with the Kimberly process.
            - At the point of sale, a consumer can input the ID of a specific diamond under consideration for purchase and access the blockchain via a well-designed user interface to learn who mined, cut, and evaluated it, complete with photos of the people involved.

      - Operating globally, it meets the SDGs of decent work and economic growth, and responsible consumption and production.
    - “These track-and-trace projects help to reduce corruption,” says Goldstein. “In some cases, they help to reduce violence, and they help to reallocate funds that are being lost in the supply chain process to people who are generally at the base of the pyramid.”
  - Document certification—Blockcerts
    - Researchers at the MIT Media Lab developed this blockchain-agnostic process to build an open standard for creating, issuing, viewing, and verifying blockchain-based certificates, with the pilot use case being university transcripts, certifications, and diplomas, enabling students/graduates to bypass the registrar’s office when needing to submit documentation to, say, an admissions office of a graduate program.
    - “It is essentially a digital wallet,” says Goldstein. “I could say I went to this school, and someone could pull up my credentials and know that I went there because it is verified and written to the blockchain that I indeed went to that school.”
- Observations and lessons learned:
  - Observations:
    - Only one-quarter of projects included in the Blockchain Impact Ledger rely on public blockchains; 29% use private blockchains, 41% use hybrid blockchains, and the remaining few percent use blockchains of unknown type.
    - Ethereum accounts for nearly half of blockchain usage, with projects leveraging smart contracts; Hyperledger captures 19%; Bitcoin, 13%; and unspecified others rounding out the list.
    - As much as Goldstein strives to evaluate the number of people served by each project, pinning down this characteristic has been difficult—so much so, that the scale of more than one-quarter of projects is unknown.
      - For projects with available data, 30% serve fewer than 1000 people, 13% serve 1000–10K people, 8% serve 10K–100K people, but somewhat encouragingly 22% serve more than 100K people.
      - Although 2017 and 2018 saw an explosion of new blockchain projects for social impact—three-quarters of projects are less than three years old—the earliest projects in the Blockchain Impact Ledger date back to 2013.
  - Lessons:
    - Scale: With so many new projects, Goldstein is not surprised that so many remain small, although the hope is that their impact will grow with time and increased exposure.
    - Failure rate: “Quite a few more social impact projects have failed than have succeeded,” she says. “Most of the projects that we have seen fail never really got off the ground.”
      - Anecdotally, failing projects have been undercapitalized, lacked sufficient buy-in, were missing necessary talent, and/or did not have appropriate product–market fit.
    - In light the struggle for projects to have the social impact they strive for, Goldstein believes it is important for any project founder to consider deeply whether blockchain is the best technology to achieve the social objective, or whether an alternative strategy might be more fitting, given the high start-up costs of a new blockchain implementation.
      - “Our sense is that a lot of the projects that failed early on tried to apply blockchain because it is sexy to a problem that doesn’t really need blockchain,” she says. “That leads to bad product-to-market fit, and it ultimately leads to a bad product.”
      - If, after careful evaluation, it is determined that blockchain can truly provide the optimal route to the project’s objective, there should be a clear understanding of the role blockchain will play—and the evaluation process should also have surfaced potential technology partners.
        - As the hype around blockchain begins to subside, its high-value applications are becoming more apparent, as well.
      - “Don’t go it alone if you don’t have to,” says Goldstein. “The ability to leverage other people’s experience is also really beneficial.”
        - Metaledger, in the pharmaceutical supply chain space and therefore not directly a social impact platform, is a consortium of pharma manufacturers and distributors that track drugs to keep counterfeit products out of the hands of consumers/patients; this is in part self-serving, but also makes a significant dent in the volume of potentially dangerous illicit drugs making their way to end users.
          - “This is much less expensive and much more efficient than if any one [pharmaceutical company] had gone it alone,” says Goldstein.

        - Although only some of the projects on the Ledger are open source, those that are can provide an instructional service to others still ascending the learning curve of blockchain implementation.
    - With innovation outpacing regulation, as it always does, some nations revert to banning the use of blockchain/cryptocurrencies; as a social impact entrepreneur or organization, it is essential to remain abreast of the prevailing rules of a target nation to avoid wasting effort and resources.
    - One recognized benefit of blockchain is disintermediation, but tamper resistance can be equally valuable.
        - Example: The anti-human-trafficking eMin project, launched by the firm Diginex partnering with the Mekong Club, uses an Ethereum-based hybrid blockchain to prevent worker exploitation by enabling workers to upload copies of their employment contracts, creating an unalterable copy of their legal rights; this serves the SDG of decent work and economic growth.
    - Many of the Blockchain Impact Ledger projects utilize private blockchains, underscoring the fact that not all transactions need be public.

2030 Blockchain for Zero Carbon and Economic and Social Resiliency—Dr. John Henry Clippinger, Swytch

- Over the years, Clippinger has applied his skills to many domains, ranging from computational law to self-sovereign data and identity, but his recent focus addresses the most pressing issue of our time: climate change and, relatedly, how to design resilient cities.
    - There is no time to waste in transitioning to a zero-carbon economy featuring socioeconomic resilience; even as we discuss this, climate models are projecting an increasingly dystopian future.
        - “We really have ten years to make fundamental, dramatic changes in how we organize our society, how we organize our economy, and how we organize ourselves in general,” he says.
    - Clippinger reiterates what we all know—changes in climate are poised to make great swaths of the planet marginally habitable at best—but he expands on these perils by adding to them the attendant financial risks, risks that could rise into the quadrillions(!) of dollars.
    - In his efforts to accelerate the necessary restabilization, Clippinger believes that blockchain-esque technologies could play a beneficial role, using *blockchain* as a stand-in for “all matters of decentralization; data-driven AI and autonomous processes; encrypted, independent, self-correcting, secure/trusted processes and certificates; peer-to-peer, open (and often open-source), digital currencies and smart tokens; and decentralized exchanges and atomic swaps—things that are in concert with nature, not in conflict with nature.”
    - The story he tells has equal parts optimism and pessimism—it all depends on which side of the take-action-now ledger you and your sphere of influence comes down on.
        - “Those countries, cities, governments, cultures, or protocols that successfully drive the 2030 transition will lead the world for the 21st century and beyond,” says Clippinger.
- Planetary risks:
    - Sea level rise and storm surge are threats to all coastal cities, with Asian metropolises at particular risk, both because of the larger number of such cities and their outsized populations.
    - Similarly, the cost of assets—buildings and infrastructure—will amount to trillions of dollars in the 2070s for individual flood-prone cities, such as Miami and New York in the West, as well as such Asian cities as Guangzhou, Kolkata, Shanghai, Mumbai, Tianjin, Tokyo, Hong Kong, and Bangkok.
        - “There will be a really dramatic financial cost,” says Clippinger.
    - The seas will be powerful, but they won’t be healthy: A World Economic Forum report anticipates that, without dramatic action now, 1M oceanic species could well be extinct by 2040.
    - Might this extinction event extend to humankind itself on a similar timeframe? Thoughtful analysis does not rule out that possibility, as desertification overtakes broad swaths of currently arable land.
- Financial risks:
    - As the planet’s life-sustaining characteristics devolve, so too will the financial climate, as expressed by Mark Carney, governor of the Bank of England, when speaking at Lloyd’s, saying “In the fullness of time, climate change will threaten financial resilience and, longer term, prosperity. The window of time to act is closing.”
        - Carney’s message reflects the need to use financial markets as a lever to turn the tide on climate change by encouraging a transition to a low-carbon economy, but that will only be possible if we base decisions on facts and act with conviction and discipline.

      - An example of the “financial mega-risk” that contributes to a total of $21T in exposure is a requirement for disclosure of fire risk by California’s insurers.
      - That is, banks are repricing real estate portfolios, readjusting margin calls, and increasing reserves in the face of projected effects of climate change.
    - In the meantime, global debt is three times the size of global stock markets, financial uncertainty is at an all-time high and has been ratcheting up over the past two decades, and the Bank of England estimates global derivative exposure to be in the $660T–$1200T range.
- When faced with this bleak geographic, institutional, and economic landscape, Clippinger sees an opportunity for wholesale structural change, transitioning from the prevailing 18th-century, top-down industrial/mechanical socioeconomic models to infrastructures that are—à la blockchain—dynamic, data-driven, decentralized, securable, and autonomous.
  - Instead of analogizing the socioeconomic system as a machine with levers, drivers, and impacts, he instead views the new modality as biological—ecological—with ebbs and flows deriving from the very interconnectedness of independent entities that compose the system.
  - Examples of this alternative viewpoint applied to system components:
    - Telecommunications—The architecture of 5G is qualitatively different than that of 4G in that service providers can build virtual end-to-end networks for particular applications (mobile broadband, machine-to-machine (IoT), high-reliability–low-latency (e.g., autonomous vehicles)).
      - With each person in the 5G network serving as not only an endpoint but also a base station, “you break up the notion of a carrier and introduce a much more decentralized system,” says Clippinger. “Once you do that, you need a very different system of governance, participation, and organization.”
        - This is the crux of his argument, yet the United States remains locked into the provider–customer model, which could put it at a disadvantage to other nations with a more forward-thinking perspective.
    - Transportation/mobility—It is not a matter of moving people between locations by traversing the two-dimensional landscape of roadways, but rather an opportunity to exploit the possibilities of movement by air (drones) or underground (subways), including the autonomous transport of objects to people, rather than the other way round.
      - The Media Lab’s City Science research group considers the notion of “cities without”: without cars, without sewers, and without other staples of the current urban landscape.
        - “Instead of sewers, we could do waste containment in a distributed way,” says Clippinger.
        - Of course, a city is more than its physical infrastructure or even the energy consumption of its fixed and mobile components; urban transformation must also consider the diversity and physical proximity of its inhabitants, their job opportunities, and residential, educational, cultural, and recreational options.
        - Intrinsic to making urban renewal a reality is devising ways to incentivize the positive transition planners envision.
          - “Some of the experiments we are running show that we can get 10x improvement in both carbon dioxide reduction and social benefit and welfare,” says Clippinger.
    - Identity/presence—The concept of a “digital twin” or 3-D overlay entails endowing a digital representation of the individual—complete with physical fidelity, social alignment, spatial uniqueness, rights, and permissions—to stand in for that person in lieu of actual physical presence.
      - More than only telepresence, Clippinger also envisions building bots on the representation.
    - Digital representation of the physical world—Like a digital twin, a spatial web of 3-D overlays of the real world enables endowing physical objects with digital identity; couple this with machine learning-based identification, and a pervasive dynamic inventory becomes possible, complete with permissions, ownership, history, and the like.
      - “There is a way of giving each object its own identity, its own wallet, its own permissions,” he says. “We’re moving not to a world of permissionless, but of micropermissions—and micropermissions that are dynamically provided and revoked and rebuilt.”
        - The firm Versus is building out this model into a what it dubs Hyperspace Transfer Protocol (HSTP), which makes possible for “parties to agree on who, what, and where everything is in physical and virtual worlds,” according to Versus’s marketing materials.
          - One example application is managing scooters in Santa Monica by sharing, bilaterally, zero-knowledge proof representations between scooters and riders; another is to aid

human or robot pickers of goods in warehouses locate and positively identify goods, eliminating error through permissioned-only picking.

    - Finance—“I like to talk about autonomous financial vehicles,” says Clippinger. “You are going to have smart contracts that will execute financial incentive mechanisms autonomously. When moving into a world that is totally distributed, you can’t have a centralized force to do that.”
    - The autonomous aspects of each such example in dynamic management rely heavily on AI, data, and unsupervised learning.
    - Energy—Clippinger’s own firm Swytch addresses dynamic electricity, both from the production and consumption points of view.
      - Swytch strives to create effective financial incentivizes for the transition from a fossil fuel-based economy into a sustainable economy.
      - Leveraging the explosion in IoT infrastructure, particularly smart meters, Swytch is working with a large bank that has installed solar panels on all its branches to track location-by-location production, compute carbon offsets accrued by each device every five minutes, and automate the generation of tradable renewable energy credits hourly, instead of having to await batched deposits of credits every two months.
        - “This makes it possible to enforce compliance [with state and local sustainability regulations], but also generates revenue,” says Clippinger.
- Assistive technologies:
  - The combination of nonfungible tokens and smart contracts—Clippinger sees this as an enabling technology to build the flexible, bottom-up, decentralized trust infrastructure necessary for positive global change.
    - Nonfungible tokens (NFTs) are, by definition, one of a kind; the combination of their scarcity and tradability confers their value.
      - Such a token provides title to an assertion made about the associated asset, whether that asset be physical, financial, solely data, a process, or a claim.
        - In fact, Clippinger anticipates a fall-off in the monetization of advertising as an overriding commercial business model in favor of tokenization.
    - Smart contracts are decentralized, open, and transparent audited processes that validate claims about an asset/entity, whether that claim is an attribute, provenance, identity, reputation, or access privilege.
      - The relying parties to the contract encode within it pertinent processes and trustworthy nonfungible tokens, whether related to privacy, permissions, reputation, identity, or otherwise.
    - Tying the pieces together is the multifunction NFT protocol, which is rooted in blockchain technology and serves as the underpinning for all manner of socioeconomically important application areas, such as energy, social media, finance, health, and mobility.
  - Stablecoins—The value of the bitcoin cryptocurrency is famously volatile, but digital currencies relegated for transactions, not speculation, serve as a stabilizing force while also conferring value.
    - Examples include the JPM Coin, Facebook’s Libra, and the Bank of England’s multipolar synthetic hegemonic currency (SHC), backed by tokens, currencies, and assets.
      - “Having one of the most respected central bankers [Carney] speaking about a whole new kind of infrastructure has a lot of implications,” says Clippinger, “not the least of which is what it will do to the dollar.”
- Transitioning the global energy infrastructure from one centralized around fossil fuels into a decentralized network of renewables will yield winners and losers—with the GDPs of the United States and Canada cumulatively losing out by as much as several trillion dollars by 2035, and Europe, India, and especially China enjoying the respective upside (in part because of its concerted national push toward refining and implementing blockchain toward its own ends, albeit in as centralized a manner as it can muster).
  - “I think we are in for a significant correction and adaptation to this new kind of order,” says Clippinger.
  - Looked at from an alternative angle, this shakeup offers an arbitrage opportunity to develop multiple, new, sustainable asset classes at the expense of traditional, stranded asset classes.
    - Trade in these new asset classes—with liquidity among them—will be conducted over blockchain networks and, predicts Clippinger, these exchanges will attract the $16T of global debt now trading in opposition to the stock market.
    - If realized according to this model, the economy could thrive even as humanity dodges the worst of climate change and establishes the basis of a sustainable future.

    - “Blockchain isn’t just a ledger,” he says. “It’s a mindset.”
- Clippinger is not naïve; he recognizes the scope of change necessary to stave off the looming climate disaster. But just as New York City transitioned from a horse-and-carriage-only transportation model in 1904 to one featuring only automobiles a scant eight years later, big change is possible when appropriate technologies and incentives are put in place and supported.

Quantum Computing: Why, How, and When?—Dr. Chris Monroe, University of Maryland and IonQ

- After decades of preparatory work, quantum computing is finally working its way out of academia and into the commercial realm, including Monroe’s own firm IonQ with its trapped-ion quantum computer.
  - Given that quantum computing sits at the confluence of quantum physics and information theory, and that the TTI/Vanguard audience is well versed in the latter, but less so in the former, Monroe sets the stage for his evaluation of the state of quantum computing first with a bit of motivational history and then with a primer on quantum mechanics.
  - Thus prepped, Monroe explains the types of problems well-suited for quantum computation.
  - Finally, he lays out the various hardware options, why he has settled on the atom as the basis for his qubit, and the status of his endeavor.
  - “Quantum computing is similar to blockchain in that there is a lot of hype behind it, but there is more than just hype and hope,” says Monroe. “There is also serious opportunity here.”
- Looking back to the dawn of the computer age, Alan Turing conceived of a universal computing machine, and Claude Shannon both conceived of the bit as the fundamental unit of computation and established the notion of information theory with the probabilistic definition of information entropy (which adopted the mathematical construct already intimately known by physicists); however, at the time, there was no hardware capable of storing bits.
  - Early computers (notably the WWII-era code-breaking Colossus (1943) and Turing-complete ENIAC (1946)) were based on physically large, fragile vacuum tubes, initiating a push toward the key hardware breakthrough of solid state electronics, with the first transistor in 1947.
    - “[That transistor] looked like a physics experiment,” says Monroe. “With it you could make logic gates, you could make memory, and so forth, but it had to have just the right amount of water, and if somebody slammed the door it didn’t work. It wasn’t very reliable; however, being solid state, it stimulated people into making it better.”
    - It took the self-fulfilling prophesy of Moore’s Law (1965) to truly launch the computer revolution, with the frequent doubling of transistor densities continuing unabated—until the perpetuation of this form of progress bumped up against the laws of physics.
      - “Transistors have gotten so small that they are now about the size of big molecules, so it is harder to squeeze more size reduction out of them,” says Monroe, “and you can’t just make the chip bigger, because then you have communication issues, latency issues, and the thing gets very hot.”
  - Even with the computer revolution in its infancy, in 1959 Richard Feynman presented *There’s Plenty of Room at the Bottom*, opining the eventual ability to construct atomic-scale circuits that would open up qualitatively new opportunities due to the properties of quantum mechanics that operate on that scale of matter.
    - “The laws of physics change when you have simple degrees of freedom, like individual atoms,” says Monroe.
    - Feynman envisioned quantum computation; Monroe and others are bringing it into being.
- The two rules of quantum mechanics (neither of which can be derived from the other, which frustrates philosophers of science):
  - Rule #1: *Quantum objects are waves and can be in superposition*.
    - The fundamental equation of quantum mechanics, Schrödinger’s equation, is a wave equation; therefore, its solution—generally denoted $|\psi\rangle$—is not localized in space, despite characterizing what might otherwise be described as a physical object.
    - “If we allow everything to be a wave, then we can talk about information as a wave,” says Monroe.
      - $|\psi\rangle = a|0\rangle + b|1\rangle$, where a and b are weighting coefficients and $|0\rangle$ and $|1\rangle$ are quantum representations of bits; $|\psi\rangle$ as a superposition specifies that, as a wave, the quantum state is an admixture of both potential states of the bit at once.
  - Rule #2: *Rule #1 only holds when you’re not looking.*

  - That is, the act of observation collapses the qubit from a superposition state $|\psi\rangle = a|0\rangle + b|1\rangle$ into an unambiguous state of either $|0\rangle$ or $|1\rangle$, with probability, respectively, of $|a|^2$ and $|b|^2$.
  - Note that in contexts where the oddities of quantum mechanics can be ignored—i.e., when dealing with macroscopic objects—we use probabilities to reflect ignorance.
    - Example: The result of a fair coin toss is described at 50/50, yet the physics of any given toss is theoretically deterministic, if one cared to consider the detailed specifics of the forces and thrust associated with it.
  - Not so with quantum mechanics: “With quantum mechanics, you *have* to use probabilities—there is no way out,” says Monroe.
- “Quantum physics isn’t hard,” he says. “It’s just weird. It only works in isolation.”
  - The simple act of measurement fundamentally changing that which is measured is, in physics, unique to quantum mechanics.
  - To consider the counterintuitive nature of quantum physics Erwin Schrödinger proposed the thought experiment of Schrödinger’s Cat, where a (decidedly macroscopic) cat is isolated in a box that that contains a flask of poison, a smaller box containing a single radioactive atom, and a Geiger counter, which is connected to a rig such that upon detecting radioactive decay it triggers a hammer to smash the flask, releasing the poison and killing the cat.
    - Prior to measurement by the Geiger counter, the atom has differential probability of having and not having decayed: it is in a superposition state.
      - This is not the paradox, because the atom is a quantum entity; the paradox of Schrödinger’s Cat is the consideration of the cat as simultaneously alive and dead, when a macroscopic, biological object can be only one or the other at any given time.
      - The “many-worlds interpretation” of quantum physics provides a workaround by considering that the universe splits into one in which the atom decays and the cat dies and one in which it does not and the cat lives.
        - “It is a neat way to wrap things up, but the problem is that there are a lot of universes, not just two,” says Monroe. “We have to consider another universe every time a quantum bit is measured, and there are gazillions of them around.”
        - This too-complicated formulation for jettisoning the second rule does not satisfy Monroe.
  - Schrödinger conjured his cat in response to a different thought experiment proposed earlier that year (1935) by Albert Einstein, Boris Podolsky, and Nathan Rosen, which attempted to blow a mathematical hole in the purported completeness—if probabilistic—of information contained in $|\psi\rangle$, the solution to the quantum mechanical wave equation.
    - This paradox involved a system of two particles (qubits), either of which have two states equally available to it; Monroe labels one qubit red and the other blue.
    - Their “entangled state” is described by the joint function $|\psi\rangle = (|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2}$, which specifies that both bits are either in state $|0\rangle$ or $|1\rangle$ with equal probability, but it is unknown which is the case pre-measurement (where the factor of $\sqrt{2}$ ensures normalization such that total probabilities sum to 100%).
    - The conundrum here is that the two systems denoted as red and blue need not be physically proximate when the measurement takes place (even if physically interacting during state preparation); that is, how is it possible for one quantum particle in an entangled system to affect the other when the pair is no longer interacting in the usual sense of the word? Wouldn’t it entail faster-than-light transfer of information for measurement of one member of the pair to instantaneously cause the other member to collapse into the correlated state?
      - Einstein called this “spooky action at a distance”; he didn’t believe it could be so—yet this is now the basis of quantum computation.
      - “If you apply information theory to this system, there is no communication of information,” says Monroe. “If you do this again and again and again—if I have many, many red ones, and each is wired to your blue ones—if I measure them, I am just going to get random zeros and ones; but I know that you have the *same* random zeros and ones, but there is zero information content in that string.”
      - No information content, but there is correlation, which is powerful for cryptography:
        - “We can have a one-time pad, without agreeing ahead of time what that pad would be,” he says. “We can wire together information without real wires.”
- Impacts for computation:

  - Parallel processing: With each qubit a superposition of two states (or more, depending on the underlying quantum entity), the number of configurations of *N* qubits grows exponentially with *N*.
    - Example: Three qubits can yield eight outputs simultaneously; more impressive is that 300 two-state qubits simultaneously yield more outputs than the number of atoms in the universe.
      - “We can compute things that we never could classically,” says Monroe.
    - “The bad news is that we can’t look at it,” he says. “When you look at it, you get only one answer, and you don’t even know which answer, due to quantum’s probabilistic nature.”
    - It wasn’t until the early 1990s that David Deutsch sorted out this decades-old riddle by recognizing the power of quantum interference.
      - “Waves can interfere, so before we make a measurement have the weightings (the set of *a*’s and *b*’s) interfere by running a circuit, where interference points guide the system,” says Monroe. “Then, before we make a measurement, there is no randomness anymore; [the output] is all weighted on one [state]. There is a forcing function on one answer, and that answer can depend on lots and lots of inputs.”
    - Quantum computation’s power derived from its ability to leverage interference, but only when a problem has many possible solutions (i.e., it has no upside for problems described by one-to-one functions).
  - Example applications that benefit from quantum computation:
    - Factoring numbers:
      - Peter Shor in 1994 developed a quantum-computing algorithm demonstrating exponential speedup, relative to classical computation.
      - Given that the basis of public key cryptography is the inability to efficiently factor large numbers, the societal implications of quantum computing are ominous.
        - The saving grace for encrypted data is the large number of qubits necessary to perform the necessary factorizations; Monroe predicts public key cryptography to remain viable for decades to come.
    - Optimization problems, even problems with a large number of variables:
      - These, by definition, have a single best solution, as well as a local neighborhood about that solution, making this class of problem well suited for quantum computation; Monroe predicts that good quantum optimizers will be in operation within a few years.
      - Examples include molecular simulation and the traveling salesman problem.
        - Len Kleinrock points out that good—and fast—heuristic solutions have been devised for selected canonical problems; it remains a matter of research whether quantum computers will be better such applications.
        - “Nobody thinks a quantum computer will be able to solve exactly something like the traveling salesman problem, but it may be able to do a better approximation than we could do classically,” says Monroe. “We’re going to find out, because it is easy to test.”
    - Big-data problems:
      - Although quantum computers are not big-data machines, “They will help us solve models that are based on big data, like the traveling salesman problem, or some logistics problem, or some allocation problem, or if you have some model of a stock market price based on thousands of inputs,” says Monroe. “It is possible that you can adjust those parameters and put them in your quantum computer to get a better approximation.”
- Quantum-computing hardware:
  - A variety of technologies are under development for use in quantum computers—none is anything like the conventional chips so familiar in classical computers.
    - These include superconducting loops, silicon quantum dots, topological qubits, diamond vacancies, neutral atoms, photonics, and—Monroe’s choice for IonQ—trapped ions.
      - The candidate technologies most likely to be successful are superconducting loops and trapped ions; he discusses the latter, which are very stable and have achieved higher gate fidelities than any of the other modalities, yet are slow and require numerous lasers.
        - The most significant downside of the use of superconducting loops is the need to operate at very low temperatures.
  - Not surprisingly, Monroe limits his hardware description to the trapped-ion model, a technology that Honeywell is also attempting to commercialize.
    - The heart of the IonQ computer is an evacuated chamber that plays host to a numerous independently levitated ions, each of which features an unpaired electron that is in a superposition of one of two states (prior to measurement).

      - Supportive structures on the chip are electrodes to trap the ions in position and lasers to perform the guidance.
      - In the earliest stages of development, the IARPA-funded trapped-ion quantum computer Monroe’s graduate students built filled an entire research lab, later iterations fit on a single tabletop, and further engineering will shrink it yet more.
      - IonQ’s current systems each support 20 qubits, which he concedes is “not yet interesting.”
      - Moving forward, the intention is to modularize the system to incorporate multicore computation, integrated photonics, a detector array, and more.
- Most advanced nations have concerted quantum-computing efforts underway, with Monroe having been involved with the crafting of the 2018 U.S. National Quantum Initiative, which authorizes $1.3B over five years and directs the National Science Foundation and National Institute of Standards and Technology to coordinate with the Department of Energy and the White House’s Office of Science and Technology Policy to move forward in this promising domain.

Spy versus Spy—Dr. Eric Haseltine, Author, and Mr. Charles L. Gandy, Retired NSA Officer and Engineer

- Highly resourced organizations have a way of overestimating their capabilities relative to those with fewer funds to throw at problems; this can lead to dangerous hubris, particularly when the organizations in question are adversarial nations, and the less prosperous one makes up for its lack of resources with clever and devious planning, and meticulous execution of those plans.
  - Such was the situation between the United States and the USSR during the Cold War, with the Soviets devising and implementing an admirable degree of tradecraft.
  - Haseltine’s book *The Spy in Moscow Station: A Counterspy’s Hunt for a Deadly Cold War Threat* tells the detailed and deeply referenced story of Gandy’s relentless quest as the head of the NSA’s R9 R&D group to root out the source of information exfiltration by the Soviets from the U.S. Embassy in Moscow—information that had led to the capture and/or deaths of an increasing number of in-country assets and CIA agents in the late 1970s.
  - This presentation to TTI/Vanguard gives only a taste of the nuanced Soviet surveillance technology and the persistence, physical discomfort, and U.S. intelligence turf wars that Gandy and his team had to endure to discover the KGB’s tactics; read the book for the full story.
    - If you wait instead for the movie to come out—yes, Haseltine and Gandy are shopping the story around Hollywood—expect alterations, both large and small, since it will be a based-on-a-true-story work of art rather the book’s rooted-in-fact, straight-from-the-horses-mouth account.
  - Haseltine and Gandy caution that U.S. overconfidence in its technological prowess is even greater today than it was back then, but so too is Russia’s ability to cobble together amazing systems to spy on us.
    - “The Russians have no money, so all they do is think—and, man, do they think!” says Haseltine. “You can’t go up against the Russians without admiring their genius. I don’t think that the same kind of espionage *could be* happening now; I know that it *is*. We wrote this book because we believe that this nation has a terrible problem being arrogant technically and not realizing what is being done to them. It was true 40 years ago; it is true today. We have a cyber-blind spot, which is about treating computers, networks, and communications systems as what they really are: not ones and zeros, but transmitters and receivers of electromagnetic energy. The Russians look at it that way. We do not look at it that way, and it is at our peril.”
- “We are going to show you a way to put a radiofrequency jammer in empty space,” says Gandy. “[With] two TV stations very loud on opposite frequencies, there is a third-order intermodulation product on the side, but there is no jamming to mask the bug. It is inside your equipment: Nonlinearities in the equipment generate the mask inside the equipment.”
  - Getting the signals out of the embassy—The antenna:
    - A special-purpose antenna discovered within a chimney adjacent to the embassy—a building with no fireplaces—was picking up the signals beneath the mask, which were transmitted by expertly engineered components added to the IBM Selectric typewriters used by embassy personnel.
    - “We looked up the chimney, and it looked like a TV antenna,” recounts Gandy, “but it wasn’t. It was three different elements for different parts of the spectrum.”
      - By entering the chimney from the next-door building, the Soviets could use cords descending from the apparatus to rotate and tilt the antenna’s alignment, or could retrieve it altogether.

      - At the time of its removal, the antenna had been focused on the office of the U.S. Ambassador, the source of the most sensitive communications.
    - Following a dramatic and hazardous covert antenna-abduction operation that Haseltine details in his book, Gandy set to analyze the apparatus using the three racks of analysis equipment that he had personally toted with him to Moscow from the United States.
      - A power supply, spectrum analyzers, receivers, variable-frequency generator, and so forth assisted him in discovering that the three elements operated at 30, 60, and 90 MHz, respectively.
      - When donning headphones, using the Soviet antenna, he surprisingly heard no trace of the two television stations broadcasting strong signals nearby, leading him to correctly assess that the control box associated with the antenna contained notch filters to eliminate the TV signals through intermodulation while simultaneously amplifying signals in the gaps between the stations.
      - The antenna elements did not pick up the stations, but signals hidden inside the intermodulation overtones that were broadcasting at intervals of 30, 60, and 90 MHz.
        - “These light switch clicks,” says Gandy, “weren’t across the spectrum, as would be a real light switch. Instead, it was in bands. That was the signal coming out.”
      - These signals would be invisible to any technical surveillance countermeasure equipment, while signals from Soviet-planted bugs/implants would come through loud and clear to the Soviet listening post.
        - Notably, such a surveillance setup would be as effective today as it was then: “This technique, which the Russians have been using for 40 years, is still undetectable today with our best bug-sweeping equipment,” says Haseltine. “You should ponder the implications of that.”
    - To demonstrate for TTI/Vanguard participants, Gandy rigged up a system in the conference hall to simulate the masking and surveillance of the Soviets.
      - “You imagine what bug-finders do,” says Haseltine. “They go around with special receivers. But they have flaws, and the Russians knew that—they knew exactly what receivers we used, and they knew that the two TV stations would generate the equivalent of an optical illusion inside our own radios, and they hid their signal underneath the phantom illusion that was in our radios.”
- Getting the communications data to exfiltrate—The typewriters:
  - Even with the antenna as proof of Soviet surveillance of the embassy, Gandy had yet to identify the source of the signals, although he had a suspicion that it was the IBM Selectric typewriters used throughout the embassy, due to a previous report by a trusted internal typewriter repairman that one otherwise damaged machine possessed an odd set of extraneous components, such as springs on switches; yet an X-ray of the machine showed no identifiable electronics.
  - Although Gandy examined other telecommunications equipment on site and was similarly stymied, he was eventually called back to Fort Meade and forbidden to pursue his Moscow-related investigation, even as information continued to leak from the embassy causing harm to U.S. personnel stationed there.
  - In 1983, the French Embassy in Moscow discovered a teleprinter with similarly curious Soviet-installed add-ons capable of transmitting stored message traffic in short bursts, akin to the “light switch clicks” Gandy had observed when in the U.S. Embassy in Moscow.
  - Finally, with go-ahead directly from the President, Gandy returned to Moscow to oversee the secure transfer of all pieces of embassy electronics to the United States for NSA’s meticulous inspection.
  - The caustic director of the project, Walt Deeley—who had bypassed orders from William Casey, the Director of Central Intelligence, and enlisted President Reagan directly to green-light the equipment recall—offered a $10K reward to the first NSA technician to find an anomaly in a piece of equipment.
    - Reagan’s edict: “Settle the issue.”
    - Ten tons of equipment—enough to fill a jumbo cargo plane—was hauled out of the embassy, as surreptitiously as possible, and transported to Fort Meade, “although the Soviets knew something was up,” says Haseltine.
  - Many thousands of man-hours and X-rays later, Mike Arenson—a young, low-level NSA tech—discovered a collection of dark circles in the X-ray of one typewriter’s bar that supports the type

ball; well-acquainted with images of normal bars, clearly something was amiss with what should be a solid piece of aluminum.

      - Further analysis of other such machines revealed the same anomaly, which turned out to be a miniscule spiral-wrapped transformer coupled to the motor to convert the output to 5 volts to power the otherwise undetected electronics within that solid bar that the Soviets had so stealthily hidden.
        - “That switch went through seven rhodium-tipped linkages to make sure it stayed connected and then went into a tiny spring on the bar that fed a current into it,” explains Gandy, who further details the use of an insulated gold-ring-backed screw to “get the antenna out and the command signals and power back into it.”
      - Gandy remains amazed by the craftsmanship of this Soviet hack: “There was no flange on it,” he says. “It was perfectly milled.”
      - Arenson indeed won the cash from Gandy and Deeley, but not full the respect that he sought from his direct supervisor.
        - In fact, Arenson lost his job because of jumping the chain of command and calling Gandy at home at 7:00 am upon his discovery.
        - Yet, all did not end badly for Arenson, who subsequently founded and later sold three companies that together netted him $500M, which he donated to charity to keep his children from acquiring an over-developed sense of entitlement.
          - “He’s a different kind of person,” says Gandy.
    - Gandy surmises that the Soviets took the opportunity to alter the West’s equipment when passing through customs, as Haseltine explains, citing a KGB defector: “When they brought [the equipment] into the airport, since they weren’t in the diplomatic pouch, they had to inspect them for customs in a special factory in Moscow. Then they ‘improved’ them and gave them back to us.”
      - Gandy adds that the equipment destined for the U.S. embassy lingered in customs for several weeks, where it was subject to modification or outright replacement.
- The intermodulation mask and altered typewriters were but two of the 18 coverup provisions implemented by the Soviets; they were nothing if not thorough.
  - Yet the technical tradecraft unearthed by Gandy is but an extension of Russian/Soviet exploits reaching back to the 1930s and 1940s, when roughly 100 microphones were discovered behind radiators in the embassy, each connecting to a dynamic microphone on building’s exterior via a 14" wooden tube.
  - The Soviets went beyond technical infiltrations, including an instance within minutes of the scheduled midnight transfer of the newly confiscated antenna to Gandy’s apartment, when a beautiful, voluptuous woman appeared at his door, asking to come in to, purportedly, retrieve clothes she had left there when the apartment had been hers in the past; in her hand was a bottle of vodka, which she invited him to share.
    - Gandy, whose apartment was full of boxes of analytical equipment, didn’t take the honeypot’s gambit, instead stepping into the hall only to find a burly thug who was surely along with the intent to do Gandy harm, although he left the confrontation unscathed.
  - The later-version technology implanted in the French Embassy’s equipment was markedly more advanced than that of the U.S. Embassy’s.
    - “We had five different generations of the 16 [implants] we found,” reports Gandy. “The [Russians] now admit there were 32. We had a generation that used batteries, and we could date the thing by how low those batteries were—at least seven years. Of these five generations, each was improved significantly over each other. Then the French came back with one—[an optical bug in a piece of OCR equipment]—that was better than any we had found.”
  - One could argue that the Soviets took advantage of the fact that elements within the U.S. government have a way of working at odds with one another.
    - Gandy was surprised when first called to Moscow by the CIA.
      - It was almost unheard of for Central Intelligence to reach out rival NSA; yet the CIA had done all that it could to sweep the agency for bugs and had come up empty-handed.
    - Throughout the years-long operation, recurring conflicts arose among the agencies as the NSA attempted to solve the security problem facing the State Department and CIA, while they, in turn, shied away from embarrassing revelations of being duped by the Soviets.
    - In the end, it took the direct Presidential intervention for Gandy to be able to finish the task he had commenced years prior.

A Survey of Blockchain Building Blocks and Techniques—Dr. Rafail Ostrovsky, University of California-Los Angeles

- The security goal of a cryptocurrency is to implement a functionality for participating parties to have a common and irreversible view of the sequence of transactions.
  - In accordance with now-conventional nomenclature, this functionality is a ledger.
  - Until recent years, such ledgers were generated and maintained by trusted third parties, such as banks, but now the notion of the distributed ledger—the blockchain—is taking hold, where there is no trusted central authority, but rather trust derived from the use of a consensus algorithm, such as proof of work (à la the Bitcoin network) or proof of stake (à la Cardano (which uses the Ouroboros algorithm) and Algorand).
  - To best understand the choices made when designing a blockchain protocol, Ostrovsky examines some of the building blocks of a blockchain network, along with their up- and downsides; along the way, he delves into a bit of the underpinning mathematics.
- The bare bones of a public transaction ledger:
  - To a first approximation, a public transaction ledger is a nonerasable whiteboard that at any moment displays the current state.
  - Participants can write to the whiteboard, but these inputs—i.e., the transactions—are filtered before being posting.
    - "The bulletin board is smart enough to check if it accepts or rejects each transaction," says Ostrovsky. "This is the functionality called *validate*. If [the submission] is correct, [the validator] appends this transaction to the bulletin board, and the bulletin board becomes bigger."
  - The order of acceptance of transactions to the state of the ledger, says Ostrovsky, "is adversarial, but not too adversarial."
    - The detailed protocols of a given blockchain specify how decisions are made, both regarding validation and the order through which transactions become part of the ledger, but the functionality of a public transaction ledger boil down to this simple set of principles.
    - As a specialized case, a trusted third party could serve as both validator and orderer, in which case "the problem is solved," he says, with a bank's internal ledger as an example application.
- Permissioned blockchain:
  - With a permissioned blockchain, a collection of somewhat-trusted servers jointly form a trusted service using multiparty computation, a mature set of decades-old technologies, which are not only easy to implement but also efficient.
    - The servers running the trusted service must be authorized—this is where the permissioning enters—but users require no authorization to submit transactions.
  - Multiparty computation:
    - Each of the n participating servers must be authorized, but not strictly trusted; collectively, they accept all submitted transactions, provided those transactions satisfy specified rules.
    - Each such server receives transactions as secret input, with server i having input x_i.
    - The goal is for the n servers to jointly compute a function over all of their inputs—$f(x_1,\ldots,x_n)$, e.g., mean or max—such that all parties learn the output function without learning the inputs of any of the other $n - 1$ servers.
      - "The magic is that you can do it without a trusted broker, just having pairwise communications and sometimes broadcast," says Ostrovsky. "You can simulate the functionality of a trusted broker without a trusted broker, and even if [some players] misbehave, they cannot subvert the computation."
      - In fact, multiparty computation provably succeeds, provided fewer than $n/3$ participants are malicious; with broadcast, that limit rises to $n/2$.
      - "These things are incredibly efficient," touts Ostrovsky. "If you have n servers and at least half can be assured to always be online and be trusted to run prescribed software, the problem is solved."
        - An obvious use case for permissioned blockchains with multiparty computation is the federated setting in which the various servers belong to different organizations, as is the case with many industrial, defense, and intelligence applications.
- Permissionless blockchain:
  - Serving as the basis for most modern cryptocurrencies, a permissionless blockchain is more complicated than its permissioned cousin.
  - Necessary building blocks:

    - Digital signatures, which consist of three distinct probabilistic polynomial-time algorithms:
      - Key generation—input is a random number; output is both a public and private key.
      - Signing—input is the private key and the message to be signed; output is a message-specific signature.
      - Verification—others can access the message-writer's public key; combining this with the encrypted message and associated signature as inputs, the output is 1 for a valid signature and 0 for an invalid signature.
      - "Correctness means that if you do everything as prescribed—you generate your public and private keys, you sign your message, and you verify your message—it should say, yes, that is a valid signature," summarizes Ostrovsky.
      - A forgery, in this construct, occurs if someone other than the owner of the private key is able to produce a signature that passes the verification test.
        - A signature scheme is *secure* if no forgeries can be generated in polynomial time.
        - A more formal specification of security—existential unforgeability—states that the scheme is robust to an adaptive chosen-message attack; that is, even if an attacker successfully dupes the owner of a public/private key pair to sign messages of the attacker's choosing, and assuming moreover that the attacker has the individual's public key, the attacker cannot go on to forge a valid signature on any message not signed by the attacker.
      - Ostrovsky notes that nothing stated thus far specifies that signatures must be unique, but there exists the stronger requirement of unique signatures, defined as being an entirely unpredictable (unforgeable) but unique random value for each distinct message that can only be revealed with the private key.
        - "Unique signatures can be used for digital lotteries," he explains. "You each have public keys. Now I announce some random numbers, and whoever has a signature that is alphabetically smallest is the winner of the lottery. Because the signatures are unique, I don't know who won. When I announce the random number, I don't know what your signature is on this random number, but once you all start announcing your signatures, we can all verify who has the alphabetically smallest signature."
        - Verifiable random functions constitute an even stronger notion yet: In this case, the signature is unique even if the public key is generated adversarially.
    - Hash functions, which serve to map an input of arbitrary length to a fixed-length value:
      - Formally, keyed hash functions are necessary, but unkeyed hash functions are used in practice; SHA256 is an example.
      - Hash functions are many-to-one functions.
    - Collision-resistance:
      - A hash function H is collision-resistant if it is practically infeasible to find distinct input values x and x' such that the hash of each is the same; that is $H(x) \neq H(x')$ for all x and x'.
    - Puzzle-friendliness:
      - For every hash output y, if a value k is sufficiently random (i.e., if it chosen from a distribution with high min-entropy), then it is hard to find an unknown value x such that the joint hash of k then x yields y; i.e., x is unfindable such that $H(k|x) = y$.
    - Random oracle:
      - A puzzle-friendly hash function H is considered a random oracle; that is, regardless the input to H, its output is random, albeit of a fixed length.
      - Bitcoin proof of work boils down to the brute-force incrementation of inputs to a random oracle, with the winner being the one to first output a hash with the requisite number of leading zeros.
    - Merkle trees:
      - A Merkle tree—or message digest—results from hashing a collection of data with a hash function, and then iteratively pairwise-hashing the output of the hash until arriving at the root of the tree.
        - The path from a leaf to the root is sufficient to prove that that leaf is a member of the tree; this is computable in log time of the number of leaves (i.e., highly efficient).
- With this set of tools, Ostrovsky returns to the core question of decentralizing a trusted ledger and considers two questions: *Who should be in charge of which transaction goes into the ledger—and when?* and *How can parties agree on the view when some of them behave maliciously?*
  - Distinction among consensus algorithms is necessary when considering these questions.
    - Proof of work:

  - Transactions to include—Each miner submits its proof of work, hoping to have it included in the blockchain; each submission is that miner’s solution *y* that solves the puzzle-friendly hash function $H(k|x)$, where *k* is now the hash of the previous block.
    - The winner of the resultant lottery (i.e., leader election) is selected according to the work invested in mining, since solving correctly for *y* (as measured by *y* possessing a specified number of leading zeroes) increases with the number of hash trials the miner computes.
  - Consensus in the face of malicious miners—A miner is defined as malicious if instead of hashing over *k* and *x*, they hash over some k' (one that does not reflect the accepted state of the blockchain) and *x*.
    - Some honest miners might adopt the adversary’s input, while others might reject it, leading to a fork in the chain—i.e., lack of consensus.
    - Another problem is that more than one honest miner might simultaneously solve the puzzle, leading to a temporal collision.
    - The solution is, when in doubt, to adopt the longest chain as the consensus outcome.
  - “The intuition is that as long as the majority of the compute power is in honest hands—and this is a formal assumption—then you can prove that the chain owned by the honest miners will grow faster than any produced by an adversary who tries to fork,” says Ostrovsky. “Therefore, there will be no forks.”
  - Although there are various formal challenges to the legitimacy of proof of work, the primary argument against it in permissionless blockchains (e.g., Bitcoin) is its inherent energy-intensivity: “The amount of power that the Bitcoin network consumes to solve these puzzles is more than the power that is being used in Ireland,” says Ostrovsky. “It is an enormous amount.”
- Proof of stake:
  - “The whole point of proof of stake is to get away from consuming so much energy to solve these enormous puzzles,” says Ostrovsky. “Instead you are just doing leader election, and if most of the players are honest, the honest players will pick what is the next block without doing the work of mining that is so damn expensive.”
  - That is, with proof of stake, instead of letting any miner submit its proof of work for inclusion in the blockchain, a lottery (leader election) is conducted from the outset.
    - “We could pick some miner at random and have this random miner decide which is the next block,” says Ostrovsky, “but instead of picking this miner at random, we will pick it according to a probabilistic method, where if you have more money in the [blockchain’s associated cryptocurrency] your chances of being picked are higher. As long as the majority of the resources—of the currency—is in good hands, the system works.”
    - This strategy is vulnerable to denial-of-service attacks if the schedule of elected leaders is known; that is, malicious participants could launch an attack on an upcoming submitter, instead substituting an adversarial block.
      - “You need to make sure that not only is this a lottery, but that this lottery is unpredictable,” he says; the use of unique signatures resolves this issue.
  - Consensus in the face of malicious miners is more challenging with proof of stake compared to proof of work, due to the ease with which an adversary can submit long fake chains.
    - The solution relies on the style of consensus originally laid out by Satoshi Nakamoto, which Ostrovsky describes as follows: “Each slot is assigned a leader with a stake-based lottery. Majority of honest stake implies that the majority of slots are allocated to honest parties (and there are different ways to do it in Ouroboros and Algorand), and we can employ the careful use of the longest chain rule to guarantee correctness—the bad guys can create small forks but not long forks—and formally what you have to prove is that the blockchain extends over time at a reasonable rate.”
      - The Ouroboros (Cardano) solution is to implement an untamperable beacon for election to ensure that progress takes place if an honest player wins the leader election.
      - The Algorand solution is to randomly select a few candidate parties according to their stake using verifiable random functions, the assumption being that a majority in this so-called slot committee will be honest (i.e., will correctly run the protocol) and therefore will successfully execute a Byzantine agreement.
        - A fresh committee is selected for each round.

- Ostrovsky concedes that this discussion is an oversimplification of the considerations that go into the construction of permissionless blockchains and, moreover, he has not addressed the matters of smart contracts or zero-knowledge-proof-based privacy due to lack of time.
- Ostrovsky and Len Kleinrock are collaborating on blockchain projects in conjunction with Kleinrock's new Connection Lab (under the auspices of UCLA's Internet Research Initiative) and the student-run Blockchain at UCLA community.

Blockchain as a Better Database—Mr. Brian Platz, Fluree

- Blockchain has emerged as a technological vehicle for the management of payments and assets, but it has the potential to do much more.
  - Specifically, Platz envisions blockchain as the enabling technology to move the majority of enterprise databases out from behind firewalls and securely put them on the public Web, where they can become transactional engines capable of ushering a new ecosystem of applications.
    - "I don't really think of blockchain and databases as two different markets," he says. "I think of them as being sort of the same thing, and I think eventually everyone else will."
  - Still, for a database to have even a minimal degree of utility, it must be able to be queried with ease; blockchains do not rise to meet this bar.
  - Platz describes how the ideas and structures of the Semantic Web can bridge the existing gap between these spaces to enable blockchain to elevate the potential of databases.
    - As he makes plain, the power of the impending network of blockchain-supported databases will be the ability to move from an "application-first model," where a single application is inextricably connected to a single database, to a model in which any given database has many consumers and any given application can access many data sources.
  - "I think blockchain will be a technology that—two years, five years, ten years from now—we will use tens or hundreds of times a day, and we won't even think or know about it, because the user experience will evolve," says Platz.
- The relationship between blockchain and a database, with payments as the connective link:
  - Purveyors of cryptocurrencies promote blockchain as a payments-as-a-platform technology; that is, payments as the end goal of a blockchain implementation.
  - Platz, in contrast, sees payments as a data state change: "If I were going to build a payment system and not use blockchain, I would be storing that data in a database. It would be transacted in an atomic transaction."
  - Generalizing, if blockchain is an effective technology for managing the data state changes associated with payments, it should be equally effective at managing changes in state of any other data type.
- Granted, there are obstacles; he addresses these, beginning by divulging the prevailing downsides of blockchain and database technologies.
  - Common implementations of blockchain (i.e., public, permissionless) are slow, bad for the environment, lack finality due to forking, and are expensive.
  - Common implementations of databases are not resistant to tampering, do not tie data to their provenance, defer data security to the consumer of the data, do not intrinsically permit data sharing (APIs overcome this), and destroy historical state by overwriting data (at best managed with log files).
  - Yet each technology can borrow from the other to overcome these deficits.
    - Permissioned blockchains can use the very consensus mechanisms used for database clusters.
      - "In theory, a blockchain can run as fast as a database can today," says Platz.
    - The cost to store a gigabyte of data on Ethereum would be millions of dollars, whereas Amazon charges just pennies; blockchain is not inherently expensive, but requires a rejiggering of priorities.
    - Blockchain, if applied to databases, would natively diffuse all the criticisms of this enterprise staple.
- The promise of the Semantic Web:
  - The first-generation, read-only Web of the 1990s and early 2000s served up information.
  - The second-generation, read–write Web commenced in 2005 with cloud computing by also serving as a client–server-based platform.
    - "This is the Web that enabled Facebook and allows us to do things like infrastructure-as-a-service," says Platz. "But this is a problem too, with the average enterprise now having over 500

        SaaS applications, most of them attempting to duplicate roughly the same data, and we're left trying to keep all of that in sync."
    - The third-generation—the machine-to-machine Web, the Semantic Web—is poised to address this fundamental deficiency as distinct data sources finally gain the ability to understand each other and intercommunicate.
        - "This is the kind of Web that we need to fulfill what we mean when we talk about digital assistants," says Platz. "How can a digital assistant—or anything—do something on your behalf if it can't figure out how to communicate with other systems."
    - Much of Tim Berners-Lee's attention over the past two decades has gone into developing the building blocks, architecture, and standards and of the Semantic Web (e.g., XML, RDF, OWL, SPARQL—all have been subjects of prior TTI/Vanguard presentations).
        - "If you ever share an article on Twitter or LinkedIn," says Platz, "and the nice little picture preview and subheader come up, that's because it is querying data from that webpage using RDF data."
    - Additional components of Berners-Lee's 2006 specification of the Semantic Web layer cake were proof, trust, unifying logic, and crypto—blockchain has the potential to tie up the loose ends of the Semantic Web that have, to date, been flapping in the wind.
        - "In order for this machine-to-machine Web to really work, we know that these things have to be solved," says Platz.
    - The Semantic Web has been incubating since the time when Edmund Muth first introduced the concept of XML to TTI/Vanguard fully two decades ago; it has endured a rough ride on the Gartner hype cycle, but it is now poised to blossom.
- To illustrate how a blockchain and database might synergistically interface, Platz presents a brief Ethereum program (a smart contract) that establishes a user's name, email, and Twitter handle as a set of strings, maps the person's Ethereum address to that user, and enables a pair of functions: the first associates those strings with the address owner's username using a *set* function; the second *gets* the values of these data (i.e., name, email, and Twitter handle).
    - "I have to know the user's ID to even get the data out of it," he says. "I can't do [a typical database call] like *select * from user* or find all the users whose name starts with A."
    - Today's workaround is to stand up an Ethereum node next to the conventional application, run it against a database to receive current data, construct an integration process to ensure that the blockchain is working only on the most up-to-date data, and finally put all of this into a traditional database to make it available for the application to query.
    - Further confounding the process is that different applications will use distinct names for the same attribute; example, one might seek *TwitterName*, another *twitter*, and a third *twitterhandle*.
- The use of resource description framework's (RDF) subject–predicate–object-structured data can unravel this mess.
    - "RDF is a beautiful data format," says Platz. "It tends to be a little verbose, but there is not a single piece of data that you cannot represent in RDF. It is the universal way of representing any sort of information."
        - Data in Wikidata (Wikipedia box) and electronic health records using HL7's FHIR's standard are all represented as RDF triples.
    - His proposition is to enforce a requirement that blockchain contracts only output state in the standardized RDF format, extended to incorporate the notion of time, which is not an elemental component of databases.
        - "Databases only understand current time," says Platz. "They don't understand historical time; instead they destroy data when you overwrite it."
    - The workaround is to append each RDF triple with a binary characterization of each change to the database: if the changing datum is being newly added, it adopts the value *true*; if it was already present and is being retracted, it adopts the value *false*.
        - Example: old state: subject = *123*; predicate = *user/email*; object = *johnny@flur.ee*; add = *false*; new state: subject = *123*; predicate = *user/email*; object = *john@flur.ee*; add = *true*
            - Within the context of the Semantic Web, an email address is part of a globally understood schema consisting of several components, including one schema associated with communication, one associated with employee contact point, and one associated with social media presence—and all of these pertain simultaneously as per contextual need.
                - "I can query across these—I can join data across these—and they don't have to sit in a single repository," says Platz.

            - Fluree’s specific approach is not to query the blockchain directly, but rather to asign the updating function to the ledger (its only output is the stream of approved and cryptographically hashed RDF changes), and then index the data (in four distinct ways, with each index a variant of a B-tree, where the leaves hold the data, chunked up by locality); each query then only pulls in the segments of the database required to answer the query.
                - This architecture delivers efficiency: “We can effectively run the database in memory,” says Platz. “This creates a database that we can, in effect, run data at the edge, across the globe, feed real-time information to it, and always have it updated.”
        - A change transaction would now occupy two rows in the database: the first (tagged *false*) removes the datum with its old value (e.g., the now-defunct email address) and one that adds it back with its new value (john@flur.ee); these two transactions would be grouped into the same block of the blockchain since they occur simultaneously.
        - It is this grouping that confers the notion of blockchain-based time to the database and makes the database immutable.
            - “Every moment in time can be reproduced,” says Platz.
        - Although this model increases the storage load associated with changing data, the cost of storage is becoming ever cheaper—a pittance compared to the value of ready access to historical data.
            - “It is not that much additional data if you can efficiently store the deltas,” he assures.
- This new database capability makes it possible for multiple, distinct organizations to jointly operate a database as equal partners, under the stipulation that all must only update data according to contextually reasonable—and clearly specified—conditions.
    - Example: An invoice may only be issued by an authorized employee of a company that, according to the ruleset, has done work for the company it is billing.
        - That is, permissions are set through transaction rules that specify the connected relationships along a defined path of references.
- “We can start building rules to secure the data we are housing alongside what is being stored and managed within the data,” says Platz.
    - With permissions built into the transaction technology, it becomes safe to publish databases.
    - “We end up with a database that is immutable and tamperproof; we end up with a database that can protect itself, and we end up with a database that has the ability to be decentralized,” he says, thus addressing all the drawbacks of databases as they have been implemented to date. “I no longer have to build APIs, because they all have real-time, verifiable updates.”
- Another powerful benefit is what Platz dubs time travel: By not overwriting database entries, queries can be conducted just as easily on historical data as on current data, including the ability to perform temporal comparisons; “Every application can instantly rewind to any moment in time, and that information is provable.”
    - Similarly, real-time applications become straightforward by tracking changes in RDF triples through time and triggering actions based on changes in state.
- With privacy-related regulations proliferating—most notably, GDPR—it becomes advantageous to not store all parties’ data in a comprehensive database, but rather to partition data to ease the process of differential removal; Platz’s scheme makes it easy to perform simultaneous queries across multiple databases *as though* they formed one database/blockchain, while enabling the wholesale deletion of a database associated with a given individual demanding to be forgotten.
    - Moreover, maintaining distance between databases, with permissions required to act on or query each individually or in combination, confers security against attack, particularly given that many recent attacks have not been against the primary databases, but rather against poorly protected copies created by contractors or partner firms with a need to do queries on the side.

Blockchain and Data-Sharing in the Automotive Industry—Dr. André Luckow, BMW Group

- Transportation has always been a cooperative endeavor by necessity: Drivers, cyclists, and pedestrians must agree on and follow a common set of rules of the road to maximize the safety of all involved, road networks must have an accepted infrastructure of traffic controls, and automakers must provide vehicles that require a common set of driving skills and include at least a minimal set of regulated safety and environmental features.
    - Yet automakers recognize that this level of coordination is only the beginning.
    - Historically, the coordination listed above was sufficient, but the rollout of AI-empowered autonomous vehicles means that it is time to up the game.

    - These days, when people conjure up images of the experience of mobility, no longer is it an image of a wind-in-the-hair road trip through a beautiful landscape; instead it is of stop-and-go traffic, honking horns, stress, and perhaps even a fender-bender.
    - This reflects poorly on the entire car-related zeitgeist, which is not in carmakers' interest.
  - But there is a solution: handing more autonomy over from drivers to the transportation network.
    - To do so safely will entail safety-critical AI systems rooted in high-quality, trustworthy, and verifiable data and algorithms.
  - "We believe that there is an immense value to using blockchain across the entire automotive ecosystem," says Luckow, "starting with the supply chain for physical and digital goods, to providing our vehicles and mobility-related services to customers, and of course for future products like autonomous driving."
  - As a starting point to achieve blockchain's potential in the mobility space, every component of the ecosystem must have a digital identity—not only individual vehicles and their components (e.g., sensors), but also external IoT and smart-city devices and customers, both drivers and riders.
  - Even this first step will require coordination and cooperation among automakers and other transportation stakeholders.
    - The Mobility Open Blockchain Initiative (MOBI) is an effort toward this end, with consortium members including several automakers, governments, and NGOs working together to devise and promote blockchain-related standards particular to this vertical.
    - A primary thrust is to develop shared data assets and decentralized data exchanges to move the industry forward.
  - While mentioning the four characteristics of future mobility that BMW recognizes—autonomous driving, connectivity, electrification, shared services (ACES)—Luckow looks to autonomous driving as the primary use case for his discussion of how the mobility experience is set to change and how blockchain and MOBI's other activities will carry the automotive sector into the future.
- Like other automakers, BMW is embracing the promise of autonomous vehicles, recognizing that not only will relief from driving improve mental health, it will also improve physical and economic health as law-abiding, cooperating vehicles suffer few accidents with attendant damage to life, limb, and property.
  - A McKinsey analysis suggests economic benefits in the United States of $800B/year from broadscale adoption of autonomous driving, while Germany anticipates $1.2B in annual healthcare savings alone.
    - Urban landscapes will also become healthier: fewer parking lots will make possible more green space, bike lanes, pedestrian walkways, and the community benefits of each of these.
  - Autonomous vehicles will not only drive on their own, but also self-manage their use and upkeep, opening up new business opportunities for automakers or other service providers in the domains of parking, charging/fueling, and payment for use.
    - "Once we reach Level-5 autonomy, we probably will not own as many cars as we do right now," says Luckow. "A lot of the consumption of mobility, very likely, will be much different, but don't ask me exactly how—it is very difficult to predict all the second-order effects of autonomous driving, but if we put mobility on the blockchain, we could enable new ways for monetization."
  - Facilitating autonomous driving and its auxiliary opportunities will rely on not only AI and IoT, but also on blockchain to provide these digital identities and secure the transactions across and among them.
    - "Everything will have an identity, and everything will be able to transact with each other," he says.
- As useful as it can be, blockchain is not a universal hammer that can pound every problem into submission.
  - Within the automotive context, criteria for use include possession of high-value data and collaboration with firms distinct from one's own.
    - "If you can do the same thing with a database, you do not need a blockchain," says Luckow.
  - Yet, these characteristics apply in the automotive realm, with blockchains and distributed consensus protocols enabling vehicles to "autonomously verify identity, enforce complex rules, negotiate with other vehicles and infrastructure, and combine behavior," according to the MOBI Grand Challenge video.
    - Goals of the Challenge are to use blockchain-related tools to achieve machine identity, position awareness, collective sensor fusion, obstacle mapping, path planning, and micropayments.
    - The intention is to go beyond decision making by independent vehicles, instead sharing information with others nearby and coordinating roadway behavior.
      - "By putting vehicles, drones, and infrastructure on blockchains, we can expand their range of effective perception and solve the congestion, safety, and environmental problems of urban

mobility," intones the video, citing the great progress made over the past 15 years since the first DARPA Grand Challenge stimulated early work in autonomous vehicles.

- With the vision laid out, next comes the hard work of achieving it; this does not happen in one fell swoop, but rather incrementally.
  - As such, an early pilot project surrounds vehicle ownership.
    - Stakeholders in the current model include the primary manufacturer (e.g., BMW), OEM firms, a bank for financing, an insurance company, and a repair shop.
      - Each has an independent view of the car and its owner, with little information sharing among them, which can introduce complications when the time comes to sell the car.
    - The target model would eliminate information asymmetry by relying on a vehicle digital passport to place the data about the car on a common blockchain to "provide a consistent view for all participants in the ecosystem of the car," says Luckow. "The owner of the vehicle is in the driver's seat regarding what data he wants to share at the end of the day, but if he wants to share, the data is always consistent."
      - The vehicle ledger would ideally include data on mileage, refueling schedule, maintenance (inspections, oil changes, diagnostics), repairs, washes, tire changes, and sale.
        - Of these, mileage is particularly pertinent, because it often serves as a proxy for vehicle condition; as such, it is the target of spoofing.
        - Inclusion in the vehicle's immutable ledger eliminates this risk; similarly with the inclusion of any accident-related repairs.
    - VerifyCar is a BMW internal prototype of the digital vehicle passport in the form of a phone app that permits access to all the real-time data of the car.
      - Among the app's features is exposure of a QR code that the owner could share with others (e.g., a bank, prospective buyer, or new repair shop) to grant them access as well.
    - The owner-facing app serves its purpose, but more generally useful is to be able to manage the digital vehicle passport according to broadly accepted standards laid out by the W3C—namely, the verifiable credential standard, which "allows us to model the relationship between the holder of a certain claim and the issuer."
      - Example: BMW, as the issuer, certifies the mileage, issues this in the form of a credential to a holder, that in turn sends it to a verifier (e.g., an insurance company or car-sharing company).
        - This verifiable credential standard does not require the raw data to change hands if, for example, a credential of acceptability is sufficient for the verifier's needs.
      - Luckow concedes that the public-key infrastructure could, on the face of it, meet the needs of this scenario, but the decentralization of blockchain technology makes dynamic information exchange possible in what could be a changeable ecosystem of stakeholders.
    - More formally, the consortium has established the MOBI Vehicle Identity as a basis to associate a vehicle with an identity that can be attached to claims from various issuers; this identity is what engages in secure transactions with others in the decentralized mobility ecosystem.
      - Components of this model are the vehicle wallet (containing the vehicle identification number and "vehicle birth certificate"—both issued by the manufacturer—and ownership certificate issued by the government, perhaps with addendum by a bank or other lender) and the vehicle data (service data, driving events, map data).
      - Access management is significantly in the owner's hands, but some aspects should be allocated, for example, to the manufacturer (e.g., safety-critical data).
        - This architecture provides an opportunity for the vehicle's owner to overtly monetize its data, perhaps by sharing map data with navigation services (instead of the implicit data collection that is currently the norm).
- Explicit use of blockchain to enhance autonomous driving:
  - On-board aspects of a self-driving car include the sensor-based perception (lidar, cameras, microphone, etc.), scene understanding, motion control, and mission/trajectory planning; offboard aspects include backend training such as image labeling, iterative retraining/improvement of models, and homologation (certification to standards/specifications/regulations) and sign-off.
    - The large parameter space makes blockchain-based data provenance a natural accelerator of this complex process.
      - "You need to secure the chain, because at the end of the day you are making safety-critical, life-or-death decisions with machine-learning models," says Luckow.
      - For instance, every sensor data point from the vehicle must be digitally signed to not corrupt evolving models or propel misdirection of the vehicle.

    - The benefits of blockchain-based data awareness, tracking, continual verification, and correctness should be clear, but how does decentralization play in?
        - In the context of autonomous driving, access to shared data benefits all in the form of better models.
            - Consider, for example, the wealth of data the Center for Advanced Transportation Technologies Lab at the University of Maryland collects from disparate sources—and the myriad ways it can slice and dice it to gain situational awareness and understand the flow of vehicles in the nation's transportation systems.
        - All benefit when multiple ecosystem partners share the data they individually collect, provided this takes place in an environment of aligned incentives.
        - Enabling technologies include blockchain, the tokenized content registry, the W3C Verifiable Credentials, crypto-incentives, and off-chain storage of data in trusted execution environments to best ensure security and privacy, along with on-chain private transactions.
            - "Crypto-incentives are key here, because you often don't know what your data is worth," says Luckow. "Why not let a cryptomarket, with shares as investment in the market, control the price of the data according to the supply and demand?"
        - "Decentralized data exchanges enable the cocreation of data, benefiting all members collectively," he says.
- Although MOBI's members recognize that the transition to autonomous vehicles will be the greatest shakeup in this industry, the benefits of blockchain are first being realized in the supply chain area.
    - Every automotive company relies on parts from an international network of suppliers, and many perform vehicle assembly in plants outside their home nation as well as within.
        - BMW's production-related blockchain pilot involves firms not only in Germany, but also in Italy, the Czech Republic, Mexico, and the United States—"and this is just the first tier," says Luckow.
    - The MOBI project aims not to only enable blockchains to provide visibility into the supply chains for each automaker individually, but for whole-ecosystem blockchains to do so throughout, with a goal to surface problems early and avoid stoppages like the one that hit North America automakers in 2018 when the Meridian Magnesium Products of America plant in Michigan—a supplier of instrument panel components to BMW, Ford, GM, and Mercedes-Benz—left each of these players without a way forward for a time.

Blockchain-Supported Federal Asylum Processing in Germany—Mr. Marcus Ziegelmeier, German Federal Office for Migration, Integration, and Asylum, and Dr. Ivan Gudymenko, T-Systems Multimedia Solutions GmbH

- As bureaucracy-heavy institutions, governments can be slow to respond to sudden changes, yet the crisis in Syria that generated a wave of refugees streaming into Europe spurred Germany's Federal Office for Migration and Refugees (Bundesamt für Migration und Flüchtlinge; BAMF) to turn to technology to ease the plight of displaced people and smooth the process of their intake and management.
    - This is not the first time that the BAMF has needed to undergo marked expansion: The office's staff quadrupled to 240 in response to the Turkish unrest that commenced in 1974; it ballooned to 4000 in the wake of the fall of the Soviet Union as asylum applications grew to 438K in 1992; had settled down to about 2500 as of 2014; and has since more than tripled once again as both Syrian civil war and ISIS-induced migration has uprooted populations en masse.
        - The peak of the recent crisis (2014–2015) saw well over 1M refugee applicants enter Germany, overrunning BAMF's resources, forcing desperate people to remain in limbo for months or even several years.
        - Systemic short-comings in 2015 included not only this staff shortage and associated processing backlog, but also the lack of identity checks, slow interagency information flow, and a deficit of workers with relevant language proficiency.
    - While personnel constitute a necessary component of its response, so too is technology, particularly given the distributed nature of BAMF itself, which while headquartered in Nuremberg is spread across dozens of offices throughout the country to meet needs in every state and region.
    - With the technical expertise of T-Systems Multimedia Solutions, the agency fast-tracked a proof-of-concept project that relies heavily on blockchain technology to modernize the operations of the BAMF, while also ensuring IT scalability with a cloud-based approach and providing basic and cross-

organizational functionality through microservices—all with a goal of better serving the people who see Germany as a prospect for a better life.

    - “We see this crisis as an opportunity to improve our office,” says Ziegelmeier, who sees the new technology infrastructure as a way to efficiently scale up/down in response to the volume of people his office must serve as global circumstances cause radical shifts in human migration.
  - Ziegelmeier focuses on the needs of his agency and its collaboration partners, while Gudymenko focuses on the technical implementation and shares a demo of the system in operation.
- The landscape BAMF must navigate:
  - While the BAMF has overarching responsibility for migration, integration, and asylum, other agencies throughout the Federal Republic interface with it, contributing to the overall mission.
    - These span municipal, state, and national-level authorities, each possessing its own siloed data, divergent legal frameworks, and diverse and sometimes-incompatible IT systems.
    - Moreover, in the age of GDPR, the need to preserve privacy looms large.
    - With top-down federal support, Ziegelmeier and Gudymenko’s tight working relationship yielded a working proof-of-concept platform in under two months—a decidedly brief development period for any project, much less an interagency endeavor using a technology not previously deployed within the government.
- Technical implementation:
  - Given that the German asylum process is complex, heterogeneous, and somewhat ill-defined, the proof of concept narrowed its scope to only consider components integral to three governmental authorities: the arrival center, which is responsible for the registration, proof of arrival, and initial distribution for every refugee; the BAMF itself, which attends to the individual’s asylum application, private hearing, notification, and final decision; and the local immigration authority, which carries out the decision, including repatriation, if deemed necessary.
    - “We narrowed down the business logic to a very specific scope,” says Gudymenko.
  - Blockchain technology lends itself well to the challenges of this project, as a BAMF video summarizes.
    - To begin, an asylum seeker’s name, country of origin, biometrics, and date of arrival enter the system via the arrivals office.
      - All agencies can securely access the record and update this single source of ground truth, while retaining the full history of transactions associated with that individual, doing so in a distributed manner, according to a given agency’s set of authorizations and need-to-know particulars.
      - For many refugees, some of this information is missing: Some lack passports for proof of origin; some have lost limbs and therefore lack fingerprints.
    - Blockchain’s native technical checks, as laid out in smart contracts, serve to support the rules and regulations of the asylum process.
    - “Information is synchronized in almost the same instant for federal, state, and municipality systems,” clarifies the video. “The federal office perceives the following aims, among others: First to ensure transparency across the entire asylum process—for all associated offices, in all steps, and from start to finish; second, to support the correct process operation in accordance with existing rules and regulations via a technical review; third, to network all involved agencies, both domestically and abroad, without the need for a central office and without the need for an agency to surrender its sovereignty.”
  - Clearly, any expectation that the various underlying authorities would suddenly jettison their legacy systems in favor of a new, unified system would be ridiculous; instead, Gudymenko has taken advantage of blockchain’s features of decentralized processing, security, and syncing, while also recognizing that the devil is in the implementation details.
    - The triad of fundamental security goals—confidentiality, integrity, availability—are difficult to simultaneously achieve, including within the blockchain paradigm.
      - Blockchains, being untamperable, do a good job regarding integrity and, due to their decentralization, inherently support availability if well designed and deployed.
      - On the other hand, permissionless blockchains compromise confidentiality due to their open readability.
        - “This will have a huge social cost, if not solved in an optimal fashion,” cautions Gudymenko, who emphasizes that the asylum proof of concept leverages a private, permissioned, proof-of-authority blockchain, “so that no one should be able to play with it before the BAMF releases it.”

        - As a privacy-preserving measure, the proof-of-concept blockchain contains only the hash associated with each asylum seeker, but not that individual’s information, per se.
          - An auxiliary field labeled *status* links out to the government’s internal ID number for the individual and n URL within the arrival center’s database to access plaintext data about the refugee.
    - “If you hear somebody say ‘Our system is secure due to the blockchain,’ that is a very dangerous, sweeping statement I often hear, and I think we should really dive into it to understand what a person really means by saying that,” cautions Gudymenko. “Usually it is a sweeping statement that doesn’t hold.”
  - In the proof-of-concept system, each participating organization retains its legacy systems, with the blockchain layer loosely coupling to those systems, thereby extending the system stack.
    - The high-level architecture retains the legacy systems of the arrival center, BAMF, and the immigration authority, respectively, as its base; atop these sits a blockchain adaptation layer, with a tailored adapter relegated to each; completing the stack is a blockchain layer in which each authority is an independent node.
      - Adjacent to the stack is a blockchain report module to probe the blockchain, on the one hand, and the legacy databases on the other, for use when evaluating the quality of the proof-of-concept output on a refugee-by-refugee basis.
    - To ensure buy-in, each agency is responsible for the entire stack pertinent to it.
      - “There is no central entity that provides and maintains this, which of course introduces a huge challenge into the system because, it being a distributed system, operations should also be carried out in a distributed fashion,” says Gudymenko. “This is a challenge we are now facing in the pilot stage, which came after the proof of concept.”
      - The eventual intention is to more tightly integrate the blockchain components into the legacy systems—or to develop new systems altogether, with blockchain technology at their core.
  - In this initial implementation, business logic is narrowed to the core steps of the asylum procedure.
  - Experimenting with both Ethereum and Hyperledger Fabric, Gudymenko settled on Fabric, appreciating the feature set of this product that was developed as a private, permissioned blockchain solution for the enterprise environment, whereas he deems Ethereum poorly suited to deployment at enterprise scale.
- Gudymenko steps through a demo of the proof-of-concept system:
  - Asylum seekers first enter the system at the arrival center, where they are issued an identification number, initiating the asylum process; there, the intake officer introduces all available information into the system—given name, surname, passport, fingerprints, arrival date, and so forth, to the extent each of these is available in a documented manner.
    - At this and any point throughout the asylum process, the system’s interface exposes permissioned components to authorized users.
      - Hashes can be displayed or hidden at the discretion of the user.
    - Thereafter, whenever an official action takes place, it is recorded into the blockchain and notice is disseminated to subscribed parties within BAMF or elsewhere in the process chain.
      - “Each time the colleague in a governmental authority does something, there is a blockchain permit that specifically defines steps, and everybody else can see at least the mere fact that there is something happening so that they can prepare themselves,” says Gudymenko.
  - The system view for BAMF workers is somewhat different: Agency-specific status labels inform them of issues and progress of the refugee.
    - “They can plan their hearings and take further actions using this tool,” he says.
  - The local immigration office can similarly track people’s movement through the system, not only enabling workers there to be prepared when directives come down from the national authority, but also issuing warnings if impending actions conflict with asylum laws or accepted procedures.
- Ziegelmeier and Gudymenko proudly note that this proof-of-concept project earned top billing at the 2019 e-Government Competition, ranking first for “best digitalization project in federal and state agencies.”
- Next steps are to integrate with additional offices and agencies in the German government for a more fully fledged pilot before attempting to expand the architecture into a multinational platform to serve the asylum-processing needs throughout Europe.

Blockchain, AI, and the GAO—Dr. Timothy Persons, U.S. Government Accountability Office

- The expressed role of the century-old Government Accountability Office is to provide auditing, evaluative, and investigative services for the U.S. Congress.
    - With science and technology inextricably tied to all facets of modern society and therefore policymaking, it should be no surprise that the GAO has assembled a dedicated Science, Technology Assessment, and Analytics (STAA) team, although the surprise might be that it did not do so sooner than this past January.
        - The intrinsic underpinnings of so many of the "wicked problems" the legislature grapples with—climate change, sustainable/affordable healthcare, human migration, to name a few—involve science; "It's not just that any one of those things is hard in and of itself," says Persons, "it's the fact that they are interconnected."
    - The largest agency in the congressional branch of government, and going by the moniker General Accounting Office until 2004, the GAO might be best known, as its prior name suggests, for its financial-auditing services, but it also has a long history of applying its investigations and evaluations to the sometimes-overhyped technical buzz phrases that swirl around.
    - As GAO's chief scientist, Persons knows it is "part of [his] job to speak to intelligent decision makers and lawmakers who are [almost universally] not scientists or engineers. They are typically shocked by the rate and convergence factors of the disruptive technologies of our day that are coming upon them, and they don't know what to do."
        - In part, this work entails providing both succinct explainer documents for the STAA's legislative client base and deep dives into particular technologies and fact-based advice on when they are/are not appropriate.
            - Complicating the matter yet further is the broad set of regulatory and oversight agencies that reach their tendrils into matters related to any given technology that shows game-changing potential.
                - For example, blockchain—even from a fintech perspective alone—is receiving attention from the Federal Reserve Board, National Credit Union, Office of the Comptroller of the Currency, the Treasury Department, the Commodity Futures Trading Commission, state banking and security regulators, as well as the alphabet soup of FDIC, BCFP, FCC, FTC, and SEC.
                - "I didn't put AI up here, because it is hard to imagine which agencies are *not* affected by AI—regulatorily, mission-wise, or otherwise," says Persons.
        - For TTI/Vanguard, he digs into two these two widely hyped technologies—blockchain and AI—sharing contexts in which he deems them useful and where the fit is less than ideal.
            - At root, the STAA provides those it serves with technology risk–benefit analyses and the information that underpins them.
- The roles of the STAA:
    - In keeping with the general mandate of the GAO to provide Congress with oversight and performance auditing, the STAA was created to fulfill this mission in areas where science and technology impact society, as well as to extend it by further providing insight and foresight; that is, to ensure continued American innovation, competitiveness, security, and well-being in a rapidly changing world.
        - "We look at a technology like blockchain and ask, 'What does this mean for us?'" says Persons. "And we give them policy options that are neutral and that are balanced, because *every* policy option has a double-edged narrative."
    - To support this work, the STAA is divided into four offices:
        - Office of Science and Technology Auditing to conduct oversight;
        - An Innovation Lab to propel governmental auditing capabilities forward by experimenting with the use of AI and/or blockchain on large federal datasets;
        - Office of Engineering Sciences to propose best practices (e.g., lifecycle cost estimation, risk analysis, agile software development) and to evaluate technology readiness;
        - Office of Technology Assessment to provide foresight and policy implications (an independent office of the same name had performed these activities for Congress during 1972–1995).
    - One set of key deliverables of the STAA are in-depth reports detailing the potential and limitations of technological innovation and science regarding society's greatest challenges.

    - Example topics: protecting the electric grid from geomagnetic disturbances, rapid response to infectious disease, scarcity of municipal fresh water, sustainability of processes and products, cybersecurity for critical infrastructure protection, and the use of biometrics for border security.
    - The STAA uses Capitol Hill caucus formation to guide the topics it explores and writes about: “When a caucus forms on the Hill on *x*, we will do a study on that,” says Persons. “Caucuses form early and often; they are not empowered the way a committee is, but they do have connectivity to committees, for instance the fintech caucus is directly connected to the House Financial Services Committee, which is where regulatory teeth come into play. We want to maximize the upsides of something like blockchain and minimize the downsides—and there are always downsides. We just need to identify them.”
  - But not everyone—including Congresspeople—have the time or will to dive deeply into a topic; for them the STAA is producing a set of single-sheet Science & Tech Spotlights; each distills out the essence of its particular emerging development with a succinct explanation along with contextually relevant opportunities and challenges.
    - The first of these Spotlights cover the topics of hypersonic weapons, probabilistic genotyping software, opioid vaccines, and blockchain and distributed ledger technologies; more are in the works.
  - The overarching goal of STAA follows Michael Hayden’s DIKW pyramid, which begins with data (D) and successively extracts more value by transforming it into information (I), knowledge (K), and ultimately wisdom (W).
    - “The key goal, when I look at our group and why we have an Innovation Lab, is to convert question into answer efficiently, and get those answers to be as wise and knowledgeable as possible,” says Persons.
- The STAA’s view of artificial intelligence:
  - Background on the state of AI:
    - The field of AI is no longer largely rooted in logical reasoning over expert knowledge encoded as rules (example: tax prep software), but is now significantly the domain of the statistical approach of machine learning that relies on model development from large training sets (example: face-recognition technology).
    - The appropriate application of this “second wave of AI”—using DARPA’s nomenclature—assumes an understanding of the consequences of probability and statistics.
    - As if this were not hard enough to implement well, the third wave will be characterized by contextual adaptation where, says Persons, “The system will do what I mean. That is much harder.”
      - An AI system that explains its reasoning to the end user will be key, but today’s digital assistants are far overcoming meeting this bar.
      - Instead it is up to users to evaluate whether the AI is following through as they intend and to nix processes if not: “There have been hype cycles over AI for many decades, but it feels more prevalent now because we can hold the Alexa device or Google Assistant in our hands,” says Persons, “but it’s still not quite like the do-what-I-mean scenario of the contextual third wave.”
      - To trust AI to important decisions independently—Level-5 autonomous driving, say, or a military system that fires when *it* deems appropriate—will require a great deal more confidence than current AIs engender.
  - High-consequence domains in which AI figures prominently—and where care must therefore be taken regarding applicability:
    - Cybersecurity, autonomous vehicles, criminal justice, and financial services.
      - Digging a little into the case of criminal justice, the goal is for AI to assist—not replace—judges as they go about their daily work of adjudicating quickly and fairly.
        - A statistical evaluation revealed an instance of systemic unfairness, not rooted in bias, but rather in judges’ hunger: On average, defendants receive lighter sentences if coming before a traffic judge early in the day rather than just before lunchtime.
        - “Judges get ‘hangry,’ too,” says Persons; knowing this enables the development of appropriate workarounds.
    - The Innovation Lab is exploring the use of machine learning across many large datasets to learn to identify improper governmental payments—i.e., fraud, waste, abuse (e.g., double payments; Medicare payments to dead people)—instead of the GAO’s current practice of rooting out improper payments in each dataset individually.

        - With improper payments expected to cost the U.S. government $1.4T over the coming decade, the effective use of AI toward this end would be a service to society.
        - The challenges will be to generate relevant training sets and to build appropriate models given the goal of simultaneously working with all transactions in the U.S. general fund—a large dataset, indeed.
- A few blockchain examples pertinent to government:
    - Department of Defense logistics:
        - For a sense of the breadth of the challenge, the DoD has almost as many distinct physical assets—nearly 5M—as Walmart has SKUs.
        - Keeping track of them all is a monumental task.
        - In contrast to Walmart, however, weapons systems tend not to be designed with obsolescence in mind; instead, they can persist for decades, with the attendant need to maintain a steady supply of replacement parts.
        - The DoD has done well in its effort to be able to attain parts when needed; what has been less successful is verifying the authenticity of those parts.
            - Persons led a 2012 covert investigation, at the joint behest of then-Senate Arms Services Committee leaders Carl Levin and John McCain, in which the GAO went on a mission to purchase replacement parts via an open Internet platform that Persons stood up.
                - The bids came in; soon too did the parts.
                    - “The quality of the counterfeits was so good that we had to take them to a non-OEM entity that did part validation by electron microscopy,” he says. “The punchline is that they were all—without exception—from Shenzhen, China, and they were on weapons system platforms like nuclear submarines and jets. There was no weapons system not touched by that.”
            - Clearly, blockchain-based supply chain management, validation, and visibility have a role to play to better secure military acquisitions.
                - “I invite all of the brainpower here to think about how we can help the Defense Department do that,” says Persons.
    - Flu shot:
        - Each year, the process for the Centers for Disease Control to determine the most appropriate strains to include, for an ample supply of vaccine to be produced, and to distribute throughout all of the nation’s cities, towns, and rural communities is another matter of supply chain coordination—hence an opportunity for blockchain technology to make a positive impact.
            - “The cycle time for our health officials is nine months,” says Persons. “It is essentially a guessing game.”
    - HIPAA:
        - It is common for people to think about the Health Insurance Portability and Accountability Act in the context of privacy, but its stated intent is for people’s health records to readily move with them as they change insurers or healthcare providers.
        - Blockchain has the potential to jointly serve the needs of privacy, portability, and information-integrity.

Real-World Blockchain Use Cases, Opportunities, and Challenges in Securing Critical Infrastructure—Dr. Michael Mylrea, Pacific Northwest National Laboratory

- “Blockchain comes with great promise, but—if we get it wrong—it comes with great peril,” says Mylrea.
    - When the domain of application is the electricity grid, that peril can translate to no lights, no heat, no communications, no traffic controls—in short, a very real threat to life.
    - For each new blockchain project he has spun up over the past two years, Mylrea first runs through a battery of questions to determine whether blockchain is relevant, and if so whether a blockchain-based solution would carry with it more risk than benefit; only if satisfied does he move forward.
        - Recognizing that any discussion of blockchain is hampered by the range of meanings people apply to the term, Mylrea defines blockchain “as a distributed database or digital ledger that records transactions of value using a cryptographic signature that is inherently resistant to modification.”
    - He shares his checklist and expounds on circumstances that might indicate this class of security solution, as well as when to tweak or forego its use.

        - Notably, Mylrea would shy away from putting highly sensitive data on a public, permissionless blockchain, given that quantum computing's eventual success in factorization will break public key encryption and leave such data open for anyone to see.
    - After laying out his collection of considerations and discussing consequences for a generalized set of applications, Mylrea looks more specifically at securing the grid, which is, in large part, an otherwise unsecured collection of IoT devices.
    - Minimally, he recommends the use of blockchain technology to track them.
        - You cannot secure an asset—or assess its risk to other assets—if you don't know you have it.
        - It is similarly detrimental to geographically confine such knowledge when that asset's failure could initiate an outage and cut off access to that very data; the distributed nature of a blockchain solves this dilemma.
- What follows is a battery of useful questions when considering blockchain for any specific application; note that these reach beyond the technology itself and to the people and processes that would use them. (In each instance, an answer of *yes* increases the likelihood that blockchain would be a good fit for the problem.):
    - Does the solution require a database?
    - Will there be multiple writers inputting/updating information?
    - Is there a lack of trust among participants?
    - Is there a lack of trusted intermediaries?
    - Can a consistent set of rules help achieve the outcome?
    - Will the governing rules be consistent over time?
    - Is transparency of the transactions an important feature?
    - Is an immutable, auditable record of transactions important?
    - Are transactions dependent or interrelated?
    - Can a distributed infrastructure reduce the risk of censorship or attack?
- Taking a good, hard look at one's problem, and making an honest assessment of blockchain's applicability to it, goes a long way toward achieving a successful outcome.
    - Blockchain is not a panacea, and should not be viewed as one.
        - Poor code, poor application, and/or poor deployment could yield a system as vulnerable as were it a centralized database.
        - Blockchain solutions tied to a specific cryptocurrency serve as ready targets for attackers and are therefore inherently vulnerable.
    - A blockchain implementation is more likely to be worthwhile when multiple writers will be inputting information and when participants lack trust among themselves or in a third party.
        - The connected digital age has ushered in a range of useful services and high-value opportunities, but the fact that enterprises—large and small—collect, aggregate, and sell our personal information has led to the general degradation of trust and privacy; blockchain sidesteps the need for trust.
        - With the degree of hype currently swirling around blockchain, it is worth assessing; surely, your competitors, customers, and partners are doing so.
- In the particular case of securing the electric grid, Mylrea has assembled a diverse team of researchers, technology, and processes at PNNL to address a range of challenges and the potential for blockchain to fill a role in their solution, particularly regarding cybersecurity optimization.
    - Although most technology solutions that seek to improve cybersecurity reduce functionality, interoperability, or data availability, blockchain does not share these downsides.
        - Instead, blockchain makes it possible to both optimize and secure a system at once.
    - The production of electricity is becoming an increasingly distributed, decentralized enterprise; managing the large collection of distributed energy resources benefits from peer-to-peer coordination.
    - Electricity is not the only resource flowing in from the edge; IoT devices are continual, timestamped sources of data about usage that can inform grid load, dynamic pricing, and even hyperlocal peer-to-peer electricity exchange, provided a system exists to support such innovations.
        - By replacing the conventional third-party intermediary (utility, electricity provider, aggregator) with the cryptographic proof of the blockchain, the grid stands to become more efficient and robust, while also enabling the growth of privacy-preserving data markets.
        - This is a welcome scenario, but it ignores the fact that the bulk of IoT devices are totally insecure, send data back to the mothership in plaintext, and rely on operating systems that have not been updated for years—if ever.

      - Mylrea has had hands-on experience breaking into IoT platforms by participating in DEFCON's IoT Village, where he was party to the tactics hackers used to exploit the vulnerabilities of all manner of off-the-shelf connected devices.
      - But an artificial environment is not the only arena for study: Ukraine has suffered two attacks on its grid.
        - Mylrea's takeaways: know/monitor assets, do not run a flat network, cybersecurity policies should protect insiders from themselves, hackers tend to use basic attacks against particularly vulnerable systems/components, and do not forgo the basics of password management control, firewalls, encryption, etc.
      - Recognizing that cyber threats evolve more quickly than defenses, he advocates for carefully crafted supply chain security and smart procurement/provisioning of devices, as well as a strategy of patching early, often, and with care; security is an ongoing process of active management.
        - Avoiding an attack is always better than executing a recovery from one.
  - The energy ecosystem is a system of things—"a system of vulnerable things," cautions Mylrea, but vulnerable things that blockchain has strong potential to protect.
    - And not only the energy ecosystem, but the entire system of resources that depend on electricity is at risk from unsecured IoT devices, which now number in tens of billions.
      - Most critical infrastructures—transportation, health, defense, education, etc., in addition to energy—comprise an intricate mix of the cyber and the physical, and therefore present an expanded attack surface to adversaries.
        - This is hardly a warning about a hypothetical risk; damage from cybercrimes is expected to reach $6T by 2021, with spending on cybersecurity on the same order of magnitude.
    - The first steps toward improvement should be to understand the IoT landscape in detail by accounting for and tracking all devices—large and small—that connect to critical systems and to securely get the most out of the data they provide.
      - "If you don't have a ledger of your critical assets, how can you know if you have been hacked?" he poses.
        - For Mylrea, the notion of IoT asset goes beyond the device to include the code base, saying, "Blockchain is exciting in that it could potentially be a secure ledger of things to better manage billions of vulnerable IoT assets and track supply chains of hundreds of millions of lines of code.
          - And for good reason: A report by cybersecurity firm CrowdStrike reports that more than half of large firms have suffered a software supply chain attack in the past year, with an average cost exceeding $1M.
      - "With blockchain you can hash the metadata, [enabling the] use of who, what, and where a transaction took place," says Mylrea. "This makes possible data provenance and granular attestation without revealing all of the sensitive information about the individual."
      - Mylrea assessed above-mentioned Ukraine grid attacks with blockchain in mind:
        - How would its use have eased recovery?
        - What data would have been kept on/off chain?
        - What would be a good choice of consensus algorithm for a grid application?
          - The recovery mechanism should not be energy intensive when the grid is compromised.
        - Other considerations include latency, throughput, interoperability, location of server nodes, and costs.
      - Beyond the realm of IoT, viable cybersecurity use cases for blockchain include management of the physical supply chain, secure records, cloud implementations, and the lifecycle of digital assets.
      - Blockchain smart contracts are well suited to generating alerts and taking action to protect the state of a machine upon the detection of an anomalous situation.
- Estonia, arguably the most digitally advanced country, is successfully applying blockchain technology across its many government platforms—e-identity, interoperability services, security/safety, healthcare, e-governance, mobility services, business/finance, and education.
  - Instead of a slow, expensive, Bitcoin-like consensus mechanism, Estonia uses one-second consensus rounds to ensure fast response for proof of participation.
  - As a small country, however, it is unclear whether Estonia's model would scale up to use across a nation the size of the United States.

    - Many elements in the long list of challenges to a blockchain implementation listed here are cultural in nature and should therefore serve as a guide for effective rollout, not inherently deter blockchain's use out of hand:
        - functionality and ease of use
        - resistance to change, culture, leadership
        - lack of legal or regulatory standards
        - workforce development and education
        - interoperability and scalability
        - making changes in immutable ledgers is tough
        - server location (including internationally)
        - transaction speed and latency
        - tie-in with legacy systems
        - human error
        - length of the blockchain and throughput
        - complex systems of systems.
    - In summary, blockchain is not a silver bullet.

Next Frontiers in Blockchains: Privacy and Interoperability—Dr. Aniket Kate, Purdue University

- When blockchain technology first hit the public's attention, it was considered a near-miracle technology that, when applied to cryptocurrencies, would endow them with the qualities of speed, low cost, decentralization, and being simultaneously digital and private.
    - The past decade has stripped some of glow from this myth; in fact, blockchain—as first laid out by Satoshi Nakamoto—is none of these things.
        - Ten transactions per second, transaction fees that can engulf the transaction itself, default recentralization due to a limited number of top-performing miners, and privacy pitfalls including de-anonymization—all of these chip away at the initial promise of blockchain.
    - Kate looks to examples from his research on people-centric privacy, and on security using cryptography and distributed computing, as he focuses on the aspects of scalability (and, relatedly, speed and interoperability) and privacy to first explain blockchain's deficiencies and then suggest trade-offs and auxiliary technologies that, when incorporated into the blockchain ethos, can deliver at least a portion of the prescribed characteristics and ancillary economic benefits promised by the hype.
    - Finally, he examines strategies for interconnecting blockchains, with the need spurred by the proliferation of blockchains and the associated cryptocurrencies that often underpin them.
        - "There is no chain to rule them all!" exclaims Kate.
- The scalability problem—and workarounds:
    - As a proxy for scalability, Kate compares transaction rates for Bitcoin (ten per second) and Visa (tens of thousands per second); three orders of magnitude is hardly incidental.
    - Some workaround solutions:
        - On-chain (layer-1) sharding—Instead of delaying the consummation of every transaction until the completion of each ten-minute Bitcoin consensus round, a transaction could complete once individual miners incorporated it into their block; this approach is gaining traction.
        - Better/faster consensus algorithms—The possibilities are many—and growing; some have been detailed elsewhere during this conference.
        - Off-chain (layer-2) payments—As a focus of Kate's research, he offers a scenario of smart contracts used to construct payment channels, which he likens to the use of a time-limited gift card.
            - Alice wishes to pay Bob in installments as he completes each phase of a multipart job.
                - Alice relegates a designated sum—say, five bitcoins—to a deposit account that is unique to the two of them.
                - As he completes each step, Bob is able to withdraw his apportioned share; if he does not complete the job in the allotted time, the remainder reverts to Alice.
            - Once the smart contract is successfully written to the blockchain, the subpayments can occur off-chain, with Bob, in the interim, keeping the latest copy of the progress toward completion.
            - "The important thing to consider here is that there are only two things that are going to the blockchain," says Kate: "the creation of such a contract, and the completion of the contract. The rest happens off-chain, just between the parties."

        - Kate notes that off-chain payments can be conducted even in blockchain implementations like Bitcoin that do not support smart contracts; those that do support them enable sophisticated off-chain possibilities.
      - If an off-chain option is possible in a given blockchain, why not create a network of blockchains, connected through payment channels?
        - Consider the following scenario: Alice wishes to transfer some bitcoin funds to Dave, but she has no direct connection to him.
          - Instead, she sends funds to Bob, who in turn directs them to Carol, who ultimately pays Dave—with each intermediary taking a sliver of the funds as payment for the service.
          - “There is a nice business there in terms of becoming an intermediary, so people will create links to such intermediaries,” says Kate. “They will provide high availability, because to perform such a transaction, they have to be online.”
          - The Lightning and Stellar networks operate in just this manner.

- The privacy problem—and workarounds:
  - With the Bitcoin protocol, the names of transaction participants do not enter the ledger, only their public keys; moreover, each user’s public key changes on a per-transaction basis, further separating identity from activity.
  - However, if Alice engages in a series of transactions using wallets with addresses (pseudonyms) A, A', A'', and A''', all such addresses appear in a linked fashion on the blockchain, which could lead to Alice’s de-anonymization if any of the various transactions becomes inadvertently published.
    - An entire subindustry of chain analysis has grown up around Bitcoin to do just this.
      - “The claim that there is good privacy with the system is just not right,” says Kate.
  - Given broad public concern about privacy—as reflected in the passage of GDPR and similar statutes in California, Canada, and elsewhere—this vulnerability of the Bitcoin network should be concerning, both to those who transact in the cryptocurrency for goods or services and to those who trade in it as an investment vehicle.
    - Similarly, other public blockchains are exposed to reputational loss, as previously enthusiastic participants worry about privacy intrusions.
  - And public blockchains are not the only problematic networks; private, consortium blockchains are also fraught with privacy issues.
    - Although members of a consortium blockchain enter into the alliance voluntarily, this does not mean that they inherently trust one another; minimally, they are marketplace competitors.
      - Some issues common to many implementations:
        - the lack of metadata access control (all participants can see all entries);
        - the immutability of blockchains does not a priori allow for the revocation of access to users who change roles within an organization or leave it altogether;
        - as is true for every system involving data, some is bound to be incorrect, yet an immutable blockchain has no provision for correction.
  - Strategies to improve blockchain privacy:
    - Forking a blockchain—e.g., Bitcoin giving way to Bitcoin Cash, or Ethereum to Ethereum Classic—undercuts the principle of immutability, but not in a structured or secure manner.
      - “We want to define provably secure ways to do these changes,” says Kate, “so that we can introduce appropriate parties with decision-making powers so they can change or remove information.”
      - Various propositions:
        - mixing, for relationship anonymity, as per CoinShuffle++, TumbleBit, and PathShuffle;
        - ring signatures, for sender anonymity, as per Monero;
        - zero-knowledge succinct noninteractive arguments of knowledge (ZK-SNARKs), as per ZeroCash;
        - confidential transactions using additive homomorphic commitments, as per ValueShuffle, Mimblewimble, and the efficient range proofs of Bulletproof.
      - “The important message,” emphasizes Kate, “is that not everything works for every blockchain scenario. You [decide on your] precise privacy requirements, and you pick something meaningful and useful in your case.”
  - With this mention of various blockchain networks, the question of connecting them naturally arises.
    - Today’s primary applications of blockchain technology are payment systems, including the cryptocurrencies themselves and auxiliary payment settlement systems, and identity and supply chain management.

  - Ultimately, Kate anticipates that the market driver for blockchain technology will be in the domain of supply chains, encompassing high-priced goods like diamonds, high-risk goods like food, highly tamperable goods like IoT devices, and versionable goods like software.
- “We need to accept the fact that these blockchains will coexist,” he says. “We have to try to understand how we will allow transactions to go through all of these different blockchains.”
  - Example uses: connect a payment blockchain to a supply chain blockchain, or connect different supply chain blockchains with each other.
- Using well-known players as stand-ins for task-specific blockchains, Kate considers a complicated transaction initiated by an operator ensconced in a traditional opaque, centralized ledger that variously engages through Ethereum, Monero, Bitcoin, and Ripple, before ultimately tapping into the resources of a Hyperledger-based permissioned blockchain.
  - This example entails the successful transfer of “something from one blockchain to another blockchain,” he says, with each intermediary perhaps extracting a fee for this service, again creating a market incentive for participation.
  - Two structures exist for each step in this process:
    - Most simple is causation, when blockchain A causes an event in blockchain B.
      - Examples include pegged sidechains, as per Cosmos, and relay chains, as per Polkadot.
    - More nuanced is dependency, where events on both blockchains depend on some additional outside event, as is the case with Interledger’s atomic swaps.
      - “We want to make sure that, only when a particular condition is satisfied, all these things could be finished simultaneously,” says Kate.
- Not surprisingly, the devil is in the details, with various interconnection challenges:
  - pairwise interoperability between each of the various protocols;
  - privacy, with the hubs at the termini of the interconnection nodes as points of vulnerability;
  - availability, with every network along an interconnected route needing to be online and activated to actualize the transaction;
  - nonresource routing, where the initiator of the transaction neither pre-knows the route through the network that the transaction will take nor trusts the intermediate networks;
  - fragmentation of transactions—à la packet switching—is not a solved problem for blockchain networks.

Securing Health Data during Analysis—Ms. Anne Kim, Secure AI Labs

- Increasingly, genetic data is proving to be the key that unlocks disease mechanisms, risk prediction, and pharmaceutical advancement—but only when a large volume of such data is available for joint analysis.
  - In particular, genome-wide association studies establish correlations of genetic information with phenotypes or diseases, leading to the discovery of predictive biomarkers.
    - “More data equals more discovery,” says Kim, who notes that the relationship between the volume of data and biomarker discovery is exponential.
  - However, concerns about medical data privacy loom large, with GDPR and HIPAA regulations codifying the need to apply serious protections to data if it is to be usable for life-saving research.
    - Adding to the challenge is that relevant data is tucked away in silos: “You have a lot of biobanks out there internationally, you have a lot of different research organizations, whether academic or industry-wide, and they are not sharing the data because of privacy and security,” she says.
  - Kim applies a suite of privacy-preserving technologies to this problem—namely, federated learning, differential privacy, secure enclaves, and blockchain—explaining the components of the problem each solves and how they work together to cover all necessary bases.
  - After detailing the technical aspects, she offers a pharmaceutical proof-of-concept case study, applying her firm’s solution to a multi-omics exploration of the gut microbial ecosystem in inflammatory bowel diseases, identifying specific bacteria types most closely associated with several different diseases—maintaining full patient anonymity with a somewhat slower computational runtime than were the analysis run in a centralized, nonprivacy-preserving manner, but still adequately quick.
- The privacy-preserving technologies:
  - Federated learning, which protects database security and data integrity.
    - Conventional learning in this space is conducted by aggregating patient data from multiple hospitals into a datalake and then performing joint analysis (e.g., two hospitals, three diseases) to

yield a principal component analysis, which indicates the extent to which the data is useful for intended comparisons.

    - This introduces two types of vulnerabilities: first during data transfer (which is a lengthy process for this volume of data), and second due to centralization.
    - The question also arises of who controls the datalake: hospital #1, hospital #2, the researcher?
  - Federated learning sidesteps these issues by retaining the raw data in each hospital’s respective database and performing computation locally: “We do a variance–covariance matrix calculation at each of these hospitals,” says Kim, with only these matrices being transferred to the researcher, where they enter into a safe distributed application (dApp) that combines them, again outputting the PCA.
    - “You can get good analysis results without ever having to move the data and without ever compromising privacy and security,” she says.
  - Still, federated learning has challenges: it has a reasonably steep learning curve, and it does not protect the algorithm that runs what is often a proprietary analysis, as revealed in the variance–covariance matrices.
- Differential privacy, which ensures true anonymization through the introduction of selective noise.
  - Typically, noise is added in the form of a Laplacian distribution, since their additive output is another Laplacian.
    - The added noise serves its purpose when the results of the relevant analysis is indistinguishable, with or without the inclusion of any given individual’s data.
    - “If I add a single sample, I should be able to protect that single sample,” says Kim. “What that means is, if you have a series of electronic health records of individuals, if you have differential privacy, you can protect every single person in that record.”
    - The operational crux of differential privacy is to add enough noise to protect the individual while also ensuring accurate results.
  - Google is easing the path for researchers by publishing its differential privacy code on GitHub and inviting its use and reuse.
    - “I encourage everyone to fork it,” says Kim.
  - The promise of differential privacy for data protection is substantial, yet as a new technology it can be difficult to implement or generalize, even with Google’s contribution.
- Secure enclaves, which protect algorithm security and database integrity.
  - This baked-into-hardware resource provides security guarantees for the physical machine that runs the various security-related software.
    - “How do you know that what that machine is doing is actually giving you correct answers?” she poses. “How do you know it hasn’t been corrupted?”
  - All modern devices pre-deploy secure enclaves in their chipsets—notably Intel SGX, ARM TrustZone, Nvidia TEE, AMD SEV—which enable remote attestation to assess both the integrity of the hardware and the protection of the code running on it.
  - This pervasive safeguard facilitates scalability by protecting the digital biomarkers or other health information resident on any device, while outperforming the alternatives of secure multiparty computation and homomorphic encryption.
  - Still, while secure enclaves protect analysis, this mode of protection is written in C++ and therefore not readily accessible in its native form to data scientists, who instead tend to have facility in R and Python.
    - “If we are asking our data scientists to also be security engineers, it is a pretty high bar for them to meet, in addition to the ones they already have to deal with in order to do their job every day,” says Kim.
- Blockchain, which maintains the consensus of operations.
  - Blockchain’s decentralization, consensus/coordination, and immutability make it possible to achieve agreement without electing a central authority.
    - Applications in the health domain include MIT’s MedRec, which is a blockchain for electronic medical records, and a collaboration among Merck, IBM, KPMG, and Walmart to track the provenance of drugs, particularly opioids.
  - Blockchain’s immutability is both a feature and a bug; “Anything you put on the blockchain you cannot change or take out,” says Kim.
    - As such, health data cannot be placed on a blockchain because of the right-of-erasure guaranteed by GDPR.

  - Summarizing, “federated learning protects data, but does not protect proprietary analysis; differential privacy protects data, but it is new and not generalizable; secure enclaves protect analysis, but it is not made for data scientists; and blockchain provides immutable consensus, but it you have to be careful with what you are putting on it because it is immutable.”
- The solution proposed by Secure AI Labs (SAIL) is to combine these four technologies, with each surmounting the shortcomings of the others:
  - By melding these into a single software solution, SAIL provides both confidential computing with secure enclaves and privacy with edge computing supported by federated learning.
  - The architecture of this approach applied to the health domain begins with the hospital encrypting its data using the secure enclave of its own server and submitting it to SAIL’s secure enclave, which receives the researcher’s analytics, operates on it, and returns the results to the researcher.
    - “Everything operates in this secure enclave such that you have the ability to keep the data where it is, technically behind the firewall of this hospital, but at the same time you can protect the analysis that is sent to the hospital and then return the results,” says Kim.
    - Blockchain enters the picture when looking under the hood of SAIL’s software:
      - The first step is to write contracts to the enclave:
        - “With every single dataset that you put in the secure enclave, you are going to put a contract of data usage,” she says; this would include terms and temporal bounds on usage.
          - This generates a means to license health data.
        - Similarly, the researcher writes a contract defining the terms for the use of the algorithm.
      - These contracts are written to the blockchain, which are submitted to a secure enclave, which enforces the contracts’ terms.
      - Post-analysis, the computational outputs export to the researcher, and the data and algorithms flush from the system.
  - To ease the process for data scientists, SAIL has crafted an interface that is simple for them to use.
    - “You just open a Jupiter notebook, and seamlessly when you do your analysis you are automatically compliant,” says Kim.
- To understand the impact of this approach, Kim makes plain the consternation that data scientists in the medical field experience all too regularly.
  - Due to the risks facing those in possession of high-value genetic data, hospitals are reticent to share granular data with researchers.
  - As such, it typically takes many months for a researcher to find a willing hospital collaborator with data about the disease the researcher is studying, endure the review of the institution’s data-sharing committee of lawyers, await the anonymization of that data, and finally receive the data itself, often through sneakernet due to the hospital’s IT team’s failure to set up a secure data transfer channel.
- SAIL’s solution:
  - Kim’s partner in the proof-of-concept project was a Boston-based pharmaceutical company interested in avoiding the conventional-process snafus by using SAIL’s secure-by-design system.
  - The scenario was to retain the relevant multi-nomic data in three respective hospitals and securely perform a federated learning analysis over those datasets while preserving the privacy of the data and the integrity of the analysis with secure enclaves.
    - A preanalysis step entailed a comparison of PCAs for federated and nonfederated learning over the datasets to ensure that those from the three hospitals would enable the joint analysis.
    - Following this was the performance of a rank-sum test to compare the identification of bacteria associated with various types of irritable bowel disease, testing results from the conventional nonfederated/nonprivate process against those from the federated/differentially private process.
      - The ranks between the two strategies were precisely the same, although the secure computation required ten minutes, compared to three minutes for the nonfederated/nonprivate computation.
  - Kim’s expectation is that the owners of health data (e.g., hospitals) will become willing partners with researchers when knowing their data will be managed with the security and privacy that the law and patients demand.
- This proof of concept involved a genome-wide association study, but it could just as well be applied to all manner of health data, incorporating behavioral data collected from wearables, health insurance claims, genomic data, clinical trials data, and so forth.
  - Moreover, analyses could be run over multiples of these, leading to innovations in personalized medicine, shared drug discovery, and more.

Internet 50: From Founders to Futurists—Dr. Len Kleinrock, TTI/Vanguard Advisory Board and UCLA

- October 29, 1969, marked the dawn of the Internet, with the first message sent—and, not incidentally, received—over a distributed network of two computers.
  - This network was the Arpanet, which, with the addition of key protocols, evolved into the Internet.
- Fifty years to the day later, Kleinrock will host a symposium celebrating this anniversary at UCLA, the institution from which Kleinrock's then-student Charlie Kline sent that message.
  - Instead of gathering a handful of people into the small lab in Boelter Hall where the work was first done, Kleinrock is inviting all comers—including members of the TTI/Vanguard community—to join him in the beautiful venue of Royce Hall for a full-day symposium exploring the reality of the current Internet and the challenges its rise and evolution present to society at large.
    - "It's about the issues—the hard problems and what needs to get fixed and addressed," he says.
- The program will encompass opening and closing keynotes, fleshed out by a potpourri of panel discussions:
  - Founding fathers of the Arpanet
  - Giant organizations of the Internet
  - The responsible Internet
  - Internet security and privacy
  - Business disruptors
  - Has true innovation stalled for decades?
  - Social disruptors
  - Future Internet technology
  - A broad vision of the future
- Speakers will include many familiar to TTI/Vanguard regulars, including Kleinrock, himself, as well as (alphabetically) Vint Cerf, Steve Crocker, Judy Estrin, Krisztina "Z" Holly, Tom Leighton, John Markoff, Bob Metcalfe, and Bud Tribble, as well as the likes of Mark Cuban, Jamie Dimon, Eric Schmidt, and Peter Thiel; Charlie Kline is also among the extended lineup.
  - "It is quite a variety," says Kleinrock—"a mix of technology people who have impacted and been impacted by the Internet itself."
- Further information and registration can be found at https://samueli.ucla.edu/internet50.