EPiC Series in Computing

Volume 63, 2019, Pages 160–169

Proceedings of 32nd International Conference on
Computer Applications in Industry and Engineering

EPiC
Computing

Elevating Beneficence in Cyberspace with Situational Trust

Kendall E. Nygard, Ahmed Bugalwi, Maryam Alruwaythi, Aakanksha Rastogi, Krishna Kambhampaty, Pratap Kotala
Department of Computer Science, North Dakota State University
Fargo, North Dakota, 58102, USA
{kendall.nygard, ahmed.bugalwi,maryam.alruwaythi,
aakanksha.rastogi,krishna.kambhampaty,pratap.kotala}@ndsu.edu

Abstract

There are myriad ways in which people benefit from systems in cyberspace that support such things as positive social interactions, electronic commerce, and automated decision making. However, harm to people and organizations can also occur, through losing privacy, fostering crime and fraud, spreading misinformation, and challenging or violating many ethical standards. Broadly characterized, systems functioning in cyberspace involve people, data, devices, computational resources, controls, and communication infrastructure. As a concept, trust refers to the state of belief in the competence of an entity to act dependably, reliably and securely within a specific situation or context. Trust is a social construct. An acceptable level of trust is essential to meaningful or satisfactory engagement and interaction among people, and, by extension, among any and all cyberspace systems. Building on the ability for entities to monitor data and drive models within contexts of how people engage when interacting with systems, we describe approaches to elevating beneficence and reducing harm and in cyberspace. We include ways in which trust is characterized and measured, relate trust and predictive analytics, and describe the potential for recent technologies like blockchains and cloud systems to help to develop a more beneficent cyberspace.

Keywords: Trust, Security, Monitoring Cloud, Purpose, Social, Blockchain

1 Introduction

The scale, capacity, and diversity of systems that function in cyberspace today are extraordinary. Computational and communication resources and systems are deep, distributed and mobile, utilizing technologies such as cloud systems, smart phones, connected devices, data sources, high-performance computing, blockchains, social media, text messaging, video conferencing, and email to carry out interactions among humans and fulfil purposes. Ideally, purposeful and beneficial outcomes are accomplished as humans engage in using available systems. Collectively, the availability of systems in cyberspace allow users to access and engage with an enormous conglomeration of people, processing, storage, and information resources. However, uniformly positive and purposeful use of such dynamic, heterogeneous, and turbulent environments is a massive challenge.

We build upon the concept of situational trust as a means of suppressing harm and supporting purposeful and beneficial outcomes. We take a wide-ranging view of trusted and purposeful use of systems, including things such as successfully enjoying entertainment, having positive social interactions, disseminating or acquiring information or knowledge, running an enterprise or business, administering health care, marketing products, or controlling systems.

Trust refers to the state of belief in the competence of an entity to act dependably, reliably and securely within a specific situation or context. Trust is a social construct. As applied to relationships and actions among people, an acceptable level of trust is essential to a meaningful or satisfactory interaction. When a person engages within a system in cyberspace, the interaction is satisfactory only if the system is trustworthy at some level.

Trustworthiness is inclusive of the usual triad of information security, consisting of confidentiality, integrity, and availability. Without secure interactions, there can be no significant level of trust. But trust goes beyond security. Lapses in cybersecurity result in bad outcomes, often through no fault of the end-user. However, bad outcomes occur in many ways that are independent of cybersecurity lapses, and underscore the importance of users continuously assessing their level of trust in the systems that they employ in seeking the ends that they pursue.

Building on the concept of trust within contexts of how people behave and engage within systems, we describe pathways for reducing harm and elevating beneficence in cyberspace. We include ways in which trust is characterized and measured, describe some ways in which malicious actors cause harm in cyberspace, and characterize some ways in which trust-based systems can potentially help to develop a more beneficent cyberspace. We tie in the role of promising newer technologies, such as blockchains, cloud computing, and orchestrated systems.

The paper is centered on the potential for trust models to be integrated into online computing and communication interactions to help head off harmful outcomes and elevate beneficence. Basing systems interactions and usage choices on trust implies that people have freedom and empowerment to manage their personal digital lives, which is not automatic in all regulatory settings.

2 Monitoring Systems and Trust

When a customer accesses a system like Amazon.com, data is gathered on a host of things, like what items are purchased, browsed, or on a wishlist; length of time lingered looking at specific products; which products have been reviewed or rated; which product reviews are examined;, and the zip code of the customer [12]. The data populate analytical models that are customized to the individual customer. Often referred to as predictive analytics, the models use methodologies like collaborative filtering, decision trees, regressions, and deep learning neural networks. As an example of customer modeling, the collaborative filtering technique basically uses metrics that compare the customer with large numbers of other customers who are similar. Identifying similar customers provides one means of making purchase recommendations for additional products by analogy with what the other customers purchased. Predictive analytics can be remarkably accurate in providing a profile of a customer in detail, in terms of things like their age, gender, income level, etc. Amazon even has a patent for anticipatory predictive analytics, providing a basis for staging or pre-shipping items to a specific warehouse or trucking locations or even the customer directly, inferring that the product will be in demand in proximity soon, even if not yet on order. Such models are effective in enhancing efficiency and maintaining high customer satisfaction. Not just limited to electronic commerce systems, predictive analytics has many applications in which there is potential to use evidence and modeling to identify the likelihood of future outcomes.

We assert that there is much commonality between predictive analytics and the concept of trust. The concept of trust has shades of meaning that involve confidence, risk, reliability, truth, belief, conviction, skepticism, and assurance. In human terms, trust refers to a degree of belief on the part of one person that another will act appropriately and with integrity within a specific context. In a fundamental trust relationship between people, in a specific context or for a certain purpose, if person A trusts person B we write A → B, where A is a trustor and B is a trustee. Broadening into entities that are not necessarily only people interacting with one another, there may be many types of contexts or situations in play. Example purposes of the interaction may be to access resources or information, acquire an opinion, control or monitor something, provide a service, or make a decision. In any case, avoiding harmful outcomes and accomplishing purposes that are in some sense beneficial purposes are basic intentions. Viewed in this way, the types of data gathering and model populating used in predictive analytics are similar to what is needed in modeling trust to help users achieve satisfactory outcomes when using systems in cyberspace.

There are also movements underway to explain and reveal biases and sensitivity to parameter values in decision-making models that can result in harm in society [6]. For example, artificial intelligence software that predicts whether or not a defendant in court is likely to reoffend is being used as a sentencing guideline, yet the methodology is unrevealed. There are many other systems in place today where transparency is lacking.

The book Weapons of Math Destruction by Cathy O'Neil concerns the societal impact of algorithms that are in use, but not well-understood [8]. Similarly, trust models are complex structures, and it is of high importance to use them carefully in influencing cyberspace interactions and personal or imposed controls.

In online systems, trusted interactions rely upon public/private key encryption, including digital signatures, but this is just one piece of the broader meaning of trust. Unacceptable outcomes often occur even when all of the communication between A and B is accurate and fully secure.

A user of a system must have a reasonably high level of trust in the system. On the flip side, and anthropomorphizing, we also assert that the system itself must have some means of trusting the user, at least for access control. Thus, there is a basic symmetry in that entities within systems must be both trusters and trustors. As an example, consider the issues surrounding an authenticated user of a system. How might a system know if a user is trustworthy even if authenticated and granted access? One approach is to monitor specific user behaviors [1]. The list below illustrates some behaviors on the part of a user that could easily be associated with untrustworthiness.

- Scanning of an important port
- Carrying a virus
- Inputting Security Sensitive keywords
- Using proxies
- Large number of unsuccessful login attempts prior to authentication
- Instantiating an unusual high or low number of logins
- Working from an unusual IP address
- Atypical time spent on the system
- Atypical frequency of usage
- Atypical usage of data storage
- Accessing the account of another user
- Atypical data error rate
- Atypical IP packet loss rate

Parameterizing these user behaviors and using them in functional expressions provide the basis for a framework for trust models that captures the relative importance of the behaviors in evaluating trust. Bayesian networks, fuzzy logic, and the analytical hierarchy process are candidate model frameworks.

Commercially available Identity and Access Management (IAM) systems have been developed. In today's digital world, IAM systems are poised to become much more than basic authentication questions, into the realm of facilitating connections that assist entry into identity-based interactions with purpose.

In 2004, Kim Cameron developed the seven laws of identity, and they have stood the test of time. In brief, the laws are as follows;

1. User Control and Consent – Information identifying a user must have consent
2. Minimal Disclosure for a Constrained Use – Limit the personal information disclosed to only what is necessary
3. Justifiable Parties – Disclosure of identifying information must be justified
4. Directed Identity – Identity systems must fulfil their purpose, yet avoid correlation handles
5. Pluralism of Operators and Technologies – Support multiple types of identify providers
6. Human Integration – The human user must be a component of the distributed system integrated via unambiguous communication between humans and machines and supported by multiple identity providers
7. Consistent Experience Across Contexts - Support separation of contexts through multiple operators and technologies

Alignment with the seven laws is critical to the instantiation of a trust-based distributed system.

Capturing and maintaining reputation data is a primary approach to employing evidence to drive trust evaluation [3,4,12]. Reputation can be defined as "an expectation about an agent's behavior based on information about or observations of its past behavior." For example, it is common in alone purchasing systems to record and publish reputation using systems such as 5-star ratings.

Reputation factors that are often used in modeling are numbers of positive and negative ratings, a history tracking of ratings, and popularity factors. Reputation is a complex concept that is not exactly the same as trust. However, reputation is often viewed as an antecedent of trust, providing a fundamental source of evidence to use in establishing trustworthiness.

3 Mathematical Modeling of Trust

A trust modeling and computational framework that can be distributed, implemented and kept current at scale is a massive task. Most trust models target P2P networks that can be seen as social networks where edges map out the relationships between peers. The edges may hold weights that are commonly employed to represent ratings. By means of performing interactions/transactions, those edges become dynamic and evolve over time. They are also relative and respective, as peers may trust or distrust one another at fluctuating and networked levels. The outcome of these model is typically a global trust score that is linked to a peer/user. This score mirrors the experiences of all peers that interacted with the holder of the trust score. The trust model described in [3] is one example. Called "TrustMe", this model is reputation-based and consists of several factors. To accommodate the many varying contexts in cyberspace, the TrustMe metric can be adapted. In systematic and formal terms, this model copes with peers/entities and interactions denoted as follows: $e \in E \equiv \{e_1, e_2, \dots, e_n\}$, and $i \in I \equiv \{i_1, i_2, \dots, i_m\}$. Entity e can be viewed as a composition of subsets of popularity, and neighbor, whereas interaction i is a composition of rating, timestamp, and context subsets.

$$e \equiv \{P \cup D\} \equiv \{\{p_1, p_2, \dots, p_n\} \cup \{d_1, d_2, \dots, d_n\}\}$$

$$i \equiv \{R \cup H \cup C\} \equiv \{\{r_1, r_2, \dots, r_m\} \cup \{h_1, h_2, \dots, h_m\} \cup \{c_1, c_2, \dots, c_m\}\}$$

$$r \equiv \{r\} \qquad \text{for one-way rating}$$

$$r \equiv \{r_{e_i}, r_{e_j}\} \qquad \text{for two-way rating}$$

These factors can be measured with a value that falls between minimum and maximum boundaries. For example, the minimum value may represent the ultimate dissatisfaction and the maximum value represents the ultimate satisfaction. The *Popularity set* represents the level of sociality of a peer/entity in a given community, and its values range between $[0, 1] \equiv [completely\ unpopular, completely\ popular]$: $0 \leq p_i \leq 1\ where\ i \in \{1, \dots, n\}$. So, trust ($T$) can be computed using the driving forces: *e* and *i*.

$$T \equiv \{t_1, t_2, \dots, t_n\}$$

$$T_e(E \cup I)$$

The trust attributes/factors collaborate together to build the trust model and generate a trust score. The following formula specifies the trust model and table 1 summarizes the factors.

$$T(e_i) = \big(\theta * d^+ * \ \big(median\,\{_{j=0}^{d^+}\, tv_j * \ p_j \ * \ h_j \ * r^+\big(e_i, e_j\big)\}\big) + (\gamma * d^- \\ * (median\,\{_{j=0}^{d^-}\, tv_j * \ p_j \ * h_j * r^-\big(e_i, e_j\big)\})))$$

Table 1. Summary of the Factors of the TrustMe Model.

Factor	Description	Range
h	indicates the prevalence of the rating historically (the age of the rating/timestamp).	[0, 1]
p	indicates the level of sociality (how the user is popular in a community).	[0, 1]
$\boldsymbol{d^+}$	denotes the centralization of the positive received ratings.	[0, 1]
$\boldsymbol{d^-}$	denotes the centralization of the negative received ratings.	[0, 1]
tv	denotes the transaction volume (context); to ignore its influence, let *tv=1.*	[$Min, $Max]
d	denotes the number of transactions that one user performs. the total number of incoming rates for a user representing the number of positive ratings and the number of negative ratings.	[0, Max]
$\boldsymbol{\theta}$	a weight for the second part of the equation; determining the influence level of this part.	[0, 1]
$\boldsymbol{\gamma}$	a weight for the second part of the equation; determining the influence level of this part.	[0, 1]

In order to guarantee the efficiency of a trust model, a fraud/deception model need to be plugged into the equation. The fraud model is a complementary model that aims to filter out dishonest feedback based on predefined criteria. Figure 1 shows a general state transition diagram of the trust model. The figure depicts how the fraud component interacts with the trust component to filter out dishonest ratings/feedback. The diagram also explains the dynamism of trust and popularity when an interaction takes place. A new trust score will be generated and updated after the ratings are filtered out by the fraud analyzer, whereas the popularity value will be updated if a new relationship is established. Additionally, the decision of categorizing a peer as trusted or distrusted is based upon the value of a threshold so that $if\ (T \geq Trust\ threshold)\ then\ state = trusted$. The diagram is divided into three sections (popularity, trust, and fraud analyzer). Each section portrays the change of states from one to another according to specified rules.

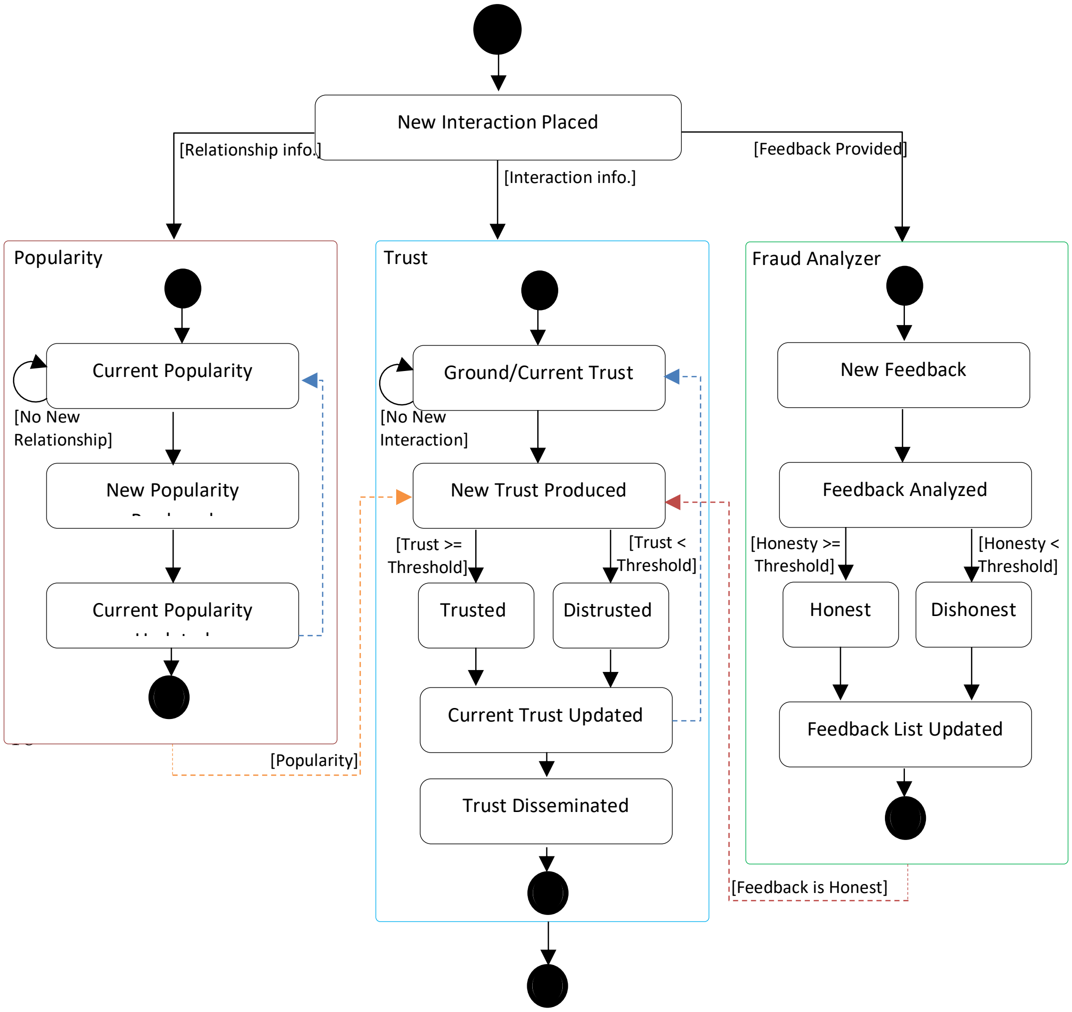

Figure 1: A State Transition Diagram of the TrustMe Model

The TrustMe model provides an indication of the computational modeling and work involved in invoking and maintaining a near real-time trust model that can be fully decentralized.

4 Situational Trust

As a social construct, trust is fundamental to how humans interact with one another. When a trustor interacts with a trustee, there are planned goals (outcomes or purposes), regardless of whether the interaction is among people or computational systems. In any case, the goals have a context, which makes them situation dependent. However, the trustee is expected to also have their own goals, which may be at cross-purposes with the goals of the trustor. Given the importance of trust in human affairs, researchers in disciplines such as sociology, psychology, economics, and political science have examined trust in various contexts and situations [11].

Usual methods for modeling trust assume a context-free environment, i.e., the trust level of an entity has no dependence on a particular situation. However, we believe that elevating the role of trust in system interactions will require situational awareness. For example, suppose that an electronic commerce transaction over the web shows that a supplier has provided a bad outcome in 3 of 10 examples in their history concerning a specific product. This poor reputation has importance in any trust model concerning that product. However, the same supplier may have a perfect reputation for another type of product, establishing situation importance. Context also pertains to the so-called “cold-start” problem, in which there is no historical record or reputation track record to provide information and guidance [3].

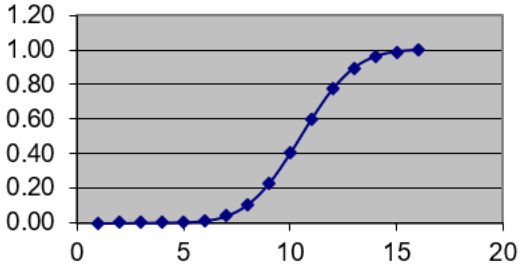

Figure 2. Hysteresis Effect in Trust

There is an old adage that we mistrust someone based upon either what we do know or what we don’t know about the trustee.

In social networks like Facebook, we normally expect truthfulness and a reasonable level of mutual trust when we interact with our friends. However, it is understood that trust levels change nonlinearly over time, rising when sequences of positive interactions take place, falling when negative interactions occur, and diminishes transitively when we interact with friends of friends or more remote systems. This is known as the hysteresis nature of trust [4]. Figure 2 illustrates the hysteresis effect in a simple model in which trust is scaled to the range [0,1], with 0 being minimum and 1 being maximum trust. Trust changes incrementally as a function of positive and negative interactions. In practice, betrayals are dramatically inimical to trust levels, while positive experiences only modestly raise them.

Cyberspace is inherently distributed, decentralized, and dynamic. The distributed nature of cyberspace means that data is generated and actions are taken at dispersed sites. Decentralized means that controls and decisions are under no fully central authority. Dynamic means that much of cyberspace is constantly changing, implying that very little can be relied upon to be static in nature. As people, devices, and resources come and go and carry out various functionalities, there is a great challenge inherent in evaluating and using trust as a basis for entering into system activity and engagement. At a minimum, there must be support for mechanisms for tracking and modeling events that are relevant to trust assessment.

5 Orchestration of Information

As system interactions proceed, more and more information that pertains to trust modeling is generated. To maintain relevant trust models, the information must be gathered, managed, and made available in prescribed ways to populate trust models that are supported. As described above, the information sources are massively distributed, decentralized, and dynamic and must be integrated in near real-time using tools produced for the purpose. Borrowing terminology from recent trends in cybersecurity, we refer the need for these sources to interoperate quickly to populate trust models as the orchestration and automation problem, or briefly, just orchestration.

In Figure 3, the Interacting Cyberspace Entities are shown as accessible systems that are sources of information, involving people, devices, cloud resources, etc. These sources are highly heterogeneous, ranging from massive NoSQL data management services, to real-time streaming of movies and music, and massively parallel processing systems.

Through real-time monitoring activity, information from diverse sources must be automatically captured and made available for orchestration within the context of the activity taking place. Automated cybersecurity-oriented orchestration would take place, but not be the only information merger supported. Context orchestration would also include tracking of the dynamic history that applies to an information source, such as a user habitually exhibiting suspicious or known bad behaviors. Multiple trust models are necessary and must themselves be merged into a combined trust measure. Trust levels associated with a source must be adjusted in nonlinear increments over time. To mitigate potential harm and to support beneficial outcomes a decision-making engine must be supported to direct action actuations that instantiate appropriate responses. Responses could range from simply continuing activity as usual, modifying how the interaction is taking place, or actually shutting down activity altogether because of low trust.

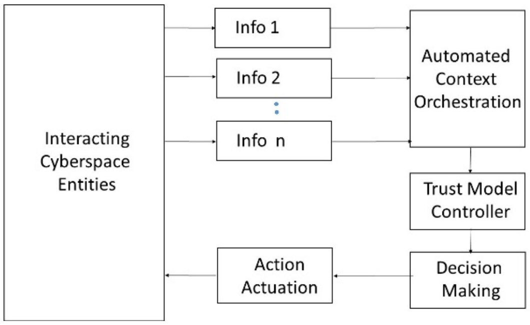

Figure 3. Trust Monitoring and Control Framework

6 Trust and Technology Changes

Blockchains are a relatively recent technological advance, with Bitcoin and other cryptocurrencies being the best-known example. Blockchains have several characteristics that result in a blockchain transaction, such as a Bitcoin operation, having a high level of trust. Blocks are stored in shared and distributed ledgers, and have a unique hash identifier that has the immutability property, in that the block cannot be changed. The first and most prominent use of blockchain technology is the capability to transfer funds transparently in a peer-to-peer fashion, avoiding the need to involve money-management organizations such as banks. By virtue of not being issued by any central authority, blockchain systems have a built-in immunity to regulation. Blockchains are quite versatile. Applications established or under development include trusted monitoring of supply chains, managing digital IDs, protecting copyrights, digital voting, transferring titles, tracking weapons, managing internet of things devices, and tracking prescription drugs.

With blockchain technology, the concept of trust is established by securing the chain using specific protocols, validating the transaction and blocks for tamper-proofing, and verifying the resources availability to guarantee the transaction execution. Essentially, blockchain systems are unhackable. Facebook has announced Libra, a dedicated blockchain-like payment system, complete with a language called MOVE for easily establishing digital contracts. The term algorithmic trust applies to blockchain systems, providing guarantees of validation and tamper-proofing of the transaction. However, algorithmic trust is not exactly the same as entity-to-entity trust.

Organizations often evaluate commitments to cloud computing carefully because of concerns for security and loss of control of their data. Hence, once an organization makes a cloud computing commitment they normally have a high level of trust that Quality of Service (QoS) agreements will be met. However, experience shows that in practice that service levels can easily fall short, quickly negatively impacting trust. This leads to the need to fully monitor performance in the cloud, including populating trust models.

Historian Niall Ferguson writes about the tension between highly distributed networked systems and hierarchical control structures [5]. The view that cyberspace respects no boundaries – geographical, political, religious, etc., is widely analyzed and acknowledged, and implies that established authorities mean little. However, history shows that systems without controls become anarchical and chaotic, and hierarchies inevitable emerge. There are many countries that have extensively censored or banned or suspended internet service within their boundaries. China, North Korea, and Iran are well-publicized examples of such countries, but there are many others. The motivations concern the idea that free speech and a free press online can result in uprisings and threaten a regime. The idea that unfettered use of cyberspace is potentially subversive and is to be feared implies that trust cannot be taken for granted.

Control hierarchies can take many forms. Facebook, Amazon, Netflix, and Google are the big four technology companies, certainly motivated by advertising, selling, and profit-making, but also hierarchical structures of great influence, strength, and power. The country of Estonia prides itself on having developed fully networked and integrated online systems, including mandatory national identity cards that are highly secure, passport management, and standardized authentication for literally all important services such as banking [10]. This extensive management of identity and digitization of all aspects of society is touted as hugely beneficial to their people, but represents a new level in government regulating, tracking, and control of citizens.

7 Conclusion

System administrators employ security software systems that carry our real-time monitoring of incoming traffic to detect and fend off malicious intruders. Defensive and offensive security procedures must be integrated to collectively manage the threats. Companies like Amazon capture real-time data and model their customers individually to serve them better and run their business efficiently and profitably. Similarly, we argue that it is feasible for arbitrary users in cyberspace to monitor and orchestrate incoming data arriving from systems that they use. These data can be orchestrated and drive trust models. These models, in turn, can support decision making to provide secure computing and satisfactory interactions and outcomes, regardless of the types of remote systems and people involved. Although challenging, the technologies exist to support a trust-based computing framework, resulting in safe, purposeful, and goal-fulfilling engagement of people and systems.

References

[1] Alruwaythi, Maryam, Krishna Kambhampaty and Kendall E. Nygard, User Behavior Trust Modeling in Cloud Security, Proceedings of the 5th Annual Conf. on Computational Science & Computational Intelligence (CSCI'18), Las Vegas, December, 2018.

[2] Abdul-Rahman, Alvarz and Stephen Hailes, "Supporting trust in virtual communities," Proceedings of the 33rd Hawaii International Conference on System Sciences, HICSS, p.6007, 2000.

[3] Bugalwi, Ahmed, Asaad Algarni and Kendall E. Nygard, A Trust Model for Bitcoin using Unsupervised Machine Learning, Proceedings of the 31st International Conference On Computer Applications In Industry And Engineering (CAINE 2018), New Orleans, 2018.

[4] Danek, Agnieszka, Joana Urbano, Ana Paula Rocha and Eugenio Oliveira, Engaging the Dynamics of Trust in Computational Trust and Reputation Systems, In: Jędrzejowicz P., Nguyen N.T., Howlet R.J., Jain L.C. (eds), Agent and Multi-Agent Systems: Technologies and Applications. KES-AMSTA 2010. Lecture Notes in Computer Science, vol 6070. Springer, Berlin, Heidelberg, 2010.

[5] Ferguson, Niall, The Square and the Tower, Penguin Press, 2018.

[6] Harris, Harlan, Transparency, Trust, and Proprietary Predictive Analytics, Medium, Feb 27, 2017.

[7] Khan, Tayyab, Karan Singh, Le Hoang Son, Mohamed Abdel-Basset, Hoang Viet Long, Satya P. Singh, and Manisha Manju, A Novel and Comprehensive Trust Estimation Clustering Based Approach for Large Scale Wireless Sensor Networks, IEEE ACCESS, May 15, .2019.

.[8] O'Neil, Cathy, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy, Crown Publishing Group New York, NY, USA, 2016.

[9] Singh, Manisha Manjul, "A Novel and Comprehensive Trust Estimation Clustering Based Approach for Large Scale Wireless Sensor Networks", Access IEEE, vol. 7, pp. 58221-58240, 2019.

[10] Reddy, Vijender Busi Atul Negi and Sashikumar Venkataraman, Trust Computation Model Using Hysteresis Curve for Wireless Sensor Networks, in IEEE SENSORS 2018.

[11] Risius, M. and K. Spohrer, K., "A Blockchain Research Framework," Bus. Inf. Syst. Eng., p. 1–6, 2017.

[12] Sterling, Bruce, Estonian e-residency, an Estonian primer, Wired, October 10, 2017.

[12] Urbano, Joana, Ana Paula Rocha and Eugénio Oliveira, A Socio-Cognitive Perspective of Trust, In Ossowski S. (ed) Agreement Technologies. Law, Governance and Technology Series, vol 8. Springer, Dordrecht, 2012.

[13] Williams, Janet, How Amazon's Focus on Data has Helped them Transform their Business, Prompt Cloud, August 10, 2018.