Cybersecurity: Balancing Efficiency with Long-Term Resilience in Connected Ecosystems

Raphael M. Reischuk

1 Connectivity as Basic Need

Many tasks of our daily life are becoming simpler due to courageous innovators and their brilliant inventions, which often depend on our intriguing new world of interconnected nodes. We can ask *Siri* how long we will need to travel to work today, we can pay our bills within seconds by scanning handy bar codes, and we can get support when choosing political candidates according to how well they represent our political views. The connectivity in our cars improves safety on the road, saves human lives, and might reduce damage to the environment. We also benefit from connected technology for tracing infection chains in a pandemic, for performing surgery at remote locations, and for assisting handicapped people throughout the day. We use digital twins of our society for finding the right partner who, hopefully, steps out of the digital network into our real life—be it for finding diligent hands in the household, for local craftsmen, or for the love of life.

In most of the above examples, the important factor that is necessary for connecting relevant stakeholders is communication. The main purpose of this type of communication, which appears fairly up in Maslow's hierarchy of needs, is to satisfy daily needs that improve the quality of life in a rather stable running society. Communication, however, is even more important during

R. M. Reischuk (✉)
Zühlke Group, Zürich, Schweiz
e-mail: Raphael.Reischuk@zuehlke.com

O. Gassmann, F. Ferrandina (eds.), *Connected Business*,
https://doi.org/10.1007/978-3-030-76897-3_13

disaster situations; it then becomes absolutely paramount. For example, one of the main concerns of aid relief is how refugees can charge their phones. Or, when Haiti was hit by a devastating earthquake in 2010, all mobile networks went down. Without connectivity, it is almost impossible for relief workers to coordinate where aid is most needed. Luckily, technology has been developed that allows mobile users to communicate directly with each other even where there is no network coverage. Another innovative solution is the mass text messaging program TERA by the *Red Cross*. It allows aid workers to send emergency messages to all smartphones in a certain region containing disaster warnings, health advice, and updates on emergency preparedness and responses.

2 A Network's Value and Cost, Intuitively

While the notion of the *value* of a network is inevitably rather vague, the underlying idea is that a network of connected nodes is becoming more valuable when there are more nodes to interact with, call, or write to. However, this notion of "build it and they will come" is defined precisely for the various kinds of networks, one of the fundamental explanations behind the oftentimes anticipated growth of a network's value is network effects.

Network effects are said to facilitate scale and are considered the main driver for successful digital businesses such as for *Apple*, *Facebook*, *Google*, *Microsoft*, *PayPal*, *Salesforce*, *Twitter*, *Uber*, and many more. The US company *NFX* attributes 70% of all value created in technology to network effects (NfX 2017) and lists 15 different kinds in its network effects map: direct network effects (physical, protocol, personal utility, personal, market network), 2-sided effects (2-sided marketplaces, 2-sided platforms, asymptotic marketplaces), data network effects, tech performance effects, and social network effects (language network effects, bandwagon effects, and belief). The less expensive it becomes to connect users on digital platforms, the faster is the growth and thus the rise of the underlying network's value. As more and more users are attracted, the existing users benefit and the network further gains competitive advantages and control due to growing market shares, which affect relationships, data, interaction, and visibility. The value of the network—may it grow exponentially, quadratically, or below following Reed, Metcalfe, or Briscoe et al., respectively—can then be understood as the ability to defend the network and to retain users.

When quantifying the value of a network, it is imperative to also consider the *costs* of running a network. It is easy to see that if all measurable costs lie within the nodes only, then clearly, the costs of the network are linear with the number of participants. But wait, is this assumption meaningful also in adversarial settings?

Metcalfe's Law

Metcalfe's law states that a network's value grows quadratically with the number of nodes, while the costs grow linearly. The intuitive explanation is that all n nodes of a network benefit from the connection to n–1 other nodes, which yields n×(n–1). The connectivity between nodes (or: participants) enables collaboration, which makes the network achieve what individuals cannot achieve on their own. The participants of a network can be computers, servers, pacemakers, industrial control panels, tablets, smartphones, and of course the human users. Simply put, a network with a single node, be it a single telegraph, a single fax machine, a single telephone, or a single user on a dating platform, has no value. With two nodes, interaction is possible and there is value—although limited—for both nodes. With 10 nodes and an assumed value of 100, when joined by 1 additional node, the network's value rises to 121. With the next node joining, the value jumps to 144 and so forth. Metcalfe's considerations, however, do not consider the vulnerability of a network as its cost: the vulnerability of a network grows more than proportionally with the size of nodes, as explained in this chapter.

3 Linear Cost also in an Adversarial Setting?

In the following, we will challenge the assumption of linear cost in light of cybersecurity. By how much is the network's value impacted when a fraction of nodes is poisoned, when malicious and careless players are inside the network? Is the spread between utility and cost in an adversarial setting as it was described by Metcalfe? If not, what are the main reasons and how can the value of a network withstand an adversarial setting?

Adversaries attempt to benefit from the prescribed network effects in the same way as legitimate users. The objectives are different, but relationships, data, and interaction are also important drivers for criminals. Some examples from different kinds of networks are as follows:

- **Platform manipulation attacks** are more effective in larger networks, simply because more targets will consume and then redistribute potentially manipulated content. Social media networks, such as *LinkedIn*, *Facebook*, and *Twitter*, in which one of the value propositions is the reliable distribution of information, rely heavily on a good reputation of not disseminating fake or manipulated content. Manipulated content can lead to mass opinion manipulation, which can undermine democracy and lead to financial and political instability.

- In fact, most platforms do suffer from targeted **spam and botnet activities**: researchers estimate that between 9% and 15% of active Twitter accounts are bots (Varol et al., 2017). Political tweets in the 2016 US election seem to originate from bots in 19% of the cases (Chaffee, 2016). For Covid-19 tweets, it appears that bots account for as much as 45% of all tweets (Allyn, 2020) and that 82% of the top 50 influential retweeters are bots, as well as 62% of the top 1000 influential retweeters (Young, 2020). Around 50 million Twitter accounts are not operated by humans (Newberg, 2017), so “likes,” “retweets,” and “followers” are totally automated and may serve a potentially criminal intention.
- Attacks based on **social engineering** require trust relationships, which are ideally established between the attacker and as many honest participants as possible. The more connections an attacker has in common with the victim, the more likely will the victim consider an interaction with the attacker legitimate. Moreover, information about the network and its participants can be leveraged by the attacker to make attacks look more authentic and less suspicious. Specific attacks based on **phishing** are best leveraged when proximity (in network parlance: good connectivity) is achievable through the network.
- In **distributed ledger networks**, it is business-critical that the majority of nodes is honest and follows the underlying protocol. The Bitcoin concept paper by Nakamoto describes the consensus mechanism as: “If a majority of CPU power is controlled by honest nodes, the honest chain will grow the fastest and outpace any competing chains” (Nakamo 2009). Consequently, if a subset of malicious nodes manages to control or influence more than 50% of all nodes, the underlying consensus mechanisms become useless. These so-called 51% attacks have been experienced a number of times recently, among others by the Ethereum-based cryptocurrencies in August 2016 and January 2019 and by Bitcoin Gold in May 2018. For cryptocurrencies, the consequences are so-called double-spending

attacks, which, in the case of Bitcoin Gold, resulted in theft of more than 18 million USD worth of Bitcoin Gold.

- In IoT networks, nodes typically rely on the **integrity of data** that is passed on to them from other nodes in the network. Reconsidering the examples mentioned above, we see the following anomalies and incidents: the digital piggy bank reports extreme savings making kids feel super rich (in case the amount is positive) or making the bank applications crash (in case the reported amount is negative). Farmers assume wrong prices on the markets and invest in the wrong products, or irrigation systems let the plants dry up. Siri takes us to wrong destinations, possibly through forbidden or dangerous areas. The bills we pay have wrong amounts and wrong payment recipients, yet they look legitimate. Political opinion making and candidate selection yields opposite results of what we would expect; however, we do not notice any anomaly. Cars crash at peak hours because traffic lights all decide to switch to green simultaneously. Contact tracing reports no infections, which renders the approach useless within days. Masses of doctors and ambulances are suddenly absorbed on site to contain the damage from failed remote surgery. Digital twins create false illusions and lead to wrong decisions and disappointments. Disaster communication reports false results and loses its availability, and therefore the trust in its utility is gone.
- In the software economy, in particular in the open-source community, value is generated by sharing source code and therefore by the trust that is placed toward **the correctness and the integrity of the source code.** As software has become the backbone of our society, novel **supply-chain attacks** appear which are deeply nested in the distribution channels of software libraries and applications. One of the most severe attacks in history is the 2020 SolarWinds attack, where the Sunburst trojan hit thousands of organizations globally, such as NATO, the EU parliament, the UK government, many US government organizations, and *Microsoft*. The cyber espionage campaign started in March 2020 with the exploitation of software and credentials from software companies such as *Microsoft*, *SolarWinds*, and *VMware*. The attackers have planted remote access tool malware into a widely used network monitoring software, which was quickly installed on tens of thousands of systems worldwide. More exploits and attacks were launched as a consequence.
- This supply-chain attack demonstrates the fragility of today's software network: a single infected piece of software can cause damage to thousands of organizations worldwide, at an incredible speed. Due to the complexity, these attacks can go undetected for months.

Put differently, the quality of the experienced network depends on the network's ability to detect and exclude malicious and otherwise unwanted participants. Striving for a high ratio of honest participants has a direct influence on the value of the network. Otherwise, unsolicited email, fake news, and the scenarios described above are a tedious and possibly enduring artifact.

Keeping unwanted attackers at bay is usually a shared responsibility between the network itself (infrastructure and operators) and the participants in the network (devices and users). These efforts clearly pay in on the cost side and must not be neglected when assessing the overall value of the network. Today's networks spend a huge body of resources on that problem and employ entire armies of digital hunters to eliminate the negative influences by human troublemakers and autonomous bots. At the same time, however, some networks invoke self-curing measures such as mutual accreditation of nodes and recommendation-based trust profiles as seen on *Uber* or *Airbnb*.

Long before Covid-19, computer and network engineers started building their architectures around the principle of "zero trust" in order to increase the overall resilience of the architectures. Essentially, this principle states that each node in the network shall assume only little honesty, correctness, and authenticity when exchanging data with other parties. In corporate environments, zero trust means companies should not generally trust their customers, employees, and applications. In particular, no differentiation is made between communication from inside and outside the company boundaries. Instead, everything and everyone who tries to access corporate data must be authenticated and controlled. It is a consistently data-centric approach based on constant monitoring. For example, in authorization decisions, access is granted only after a consequent and reliable identification of the requesting party. As such, the trust boundaries vanish at the expense of a growing cost for scrutinizing devices, humans, and network traffic. Applied correctly, these costs are incurred at every node of a network and thus scale linearly with the number of nodes and data traffic exchanged between the nodes.

Turning back to Metcalfe and Reed: if quadratic or exponential laws were true, there would be overwhelming incentives for all networks of similar technology to merge or at least to interconnect. There is, however, only little evidence for such mergers: in the late 1980s and early 1990s, the commercial online companies such as *CompuServe*, *Prodigy*, *AOL*, and *MCI Mail* provided e-mail services to subscribers, but only within their own systems.

Complete networking was not achieved until the mid-1990s. Why has it taken so long? Shouldn't the incentives by the laws of Metcalfe and Reed make isolated networks totally obsolete? The answer is no; still today, plenty of isolated networks exist. A prominent example for non-interconnected networks in the times of Covid-19 is the various national contact-tracing systems, which are not interoperable, mainly due to non-technical reasons.

Isolated networks such as VPNs, closed-user groups, and zero-trust zones serve also a dedicated purpose, inter alia for cybersecurity reasons: isolated networks usually have a high entry barrier, come with minimal trust bases, and offer additional amenities such as free expression of opinion, non-traceability, anonymity, and confidentiality guarantees.

Summarizing the above, in particular when considering aspects of cybersecurity, the value of a network seems far from growing quadratic or exponential. The question we will try to answer in the following is: what does it need to bring the value of connected and distributed systems back to growing beyond linear with the number of nodes?

4 Cyber Resilience Strategies to the Rescue

One of the main differentiating factors compared to early computer networks in the 1970s is the agility and dynamics of today's networks. Hyperscalers spin up containers, virtual machines, and networks in a matter of milliseconds. Messages, posts, media, and other data objects propagate through the nodes of a network with largely varying speeds and densities. Today's software release cycles allow for delivering new functionality and patches to billions of devices within minutes. At the same time, attackers orchestrate entire armies of bots within a blink of an eye and distribute malware to billions of devices at the same pace.

Consequently, protecting the nodes of a network (and therewith preserving the network's value) requires adaptive concepts that are of similar dynamics and agility as today's networks. Any protection concept must adapt quickly to modified circumstances inside the network. Additionally, not only changes of internal network conditions must be considered, but also external threats and attacks that are highly dynamic and often unpredictable must be incorporated in the mitigation concept.

Static approaches such as fixed firewall rules, static network zones, and fixed permission models are no longer adequate. A meaningful analogy to those static measures is a solid rock in the surf. It has a ponderous and unwieldy appearance and would never adapt to changing circumstances. A displacement of the rock by even one millimeter seems infeasible. In other words: no agility.

Cyber resilience is the exact opposite of that rock. Resilience in general describes the ability of a system, an organization, or a network to recover readily from adversity, illness, anomaly, and other unwanted condition. It is the capacity to adapt to disruptive changes in the environment, be it from adversarial input, manipulation, or abuse. Resilient systems are typically characterized by the two features: diversity and redundancy.

A meaningful analogy to illustrate the concept of resilience is a bamboo forest. Bamboo trees have distinguished properties: they are evergreen, they belong to the fastest-growing plants in the world due to their unique rhizome-dependent system (certain species grow up to 1 meter within 24 h), they have a high strength-to-weight ratio making them useful for building structures and withstanding external influence, and they are a source of food and generally constitute a versatile raw product. Being fast-growing and withstanding external changes and being versatile and appearing in large numbers correspond to diversity and redundancy and are thus important key factors for surviving disruptive changes in the environment.

How does the concept of resilience translate to networks and their nodes? As mentioned earlier, most networks carry within them a high degree of dynamics. Nodes come and join, leave, and come back. Networks are very versatile and often unpredictable. This is why most nodes in a network will likely not know all other nodes. In particular, a node might not be able to assess the trustworthiness of other nodes. Nodes can sometimes not even be certain about their direct neighborhood. In order to survive and not become prey of malicious nodes in the network, one promising concept as part of the resilience strategy is the aforementioned concept of zero trust. Despite not always being applicable and sometimes being reduced to “little trust,” the concept seeks at benefitting from network effects, but not suffers from typical trust weaknesses as mentioned in the examples above.

Robust Measures for More Network Resilience

Continuous monitoring and anomaly detection: The "zero-trust" paradigm requires probing and scrutinization of all relevant objects on the basis of which business-critical decisions are made. Data sources that are typically fed into machine-learning-based solutions include network traffic between devices, system and application log files, user behavior, dedicated honey pots, and other forms of statistical data.

Incident response and digital forensics: The "assume breach" paradigm, as the name suggests, assumes system compromise and an attacker who is closer and more successful than we think. Consequently, defense expects an incident at any moment and is prepared for various scenarios. Since 100% security is not even true in our dreams, incident response plans should be worked out, trained, and reviewed. A digital forensic team for identifying the root cause and for containing the attack should be quickly accessible upon request.

Regular penetration testing of infrastructure and applications: When constant monitoring of an infrastructure or an application is not possible at runtime, periodic testing of all critical systems should be conducted. The findings should be presented to the executive board and should be addressed with sufficient resources, diligence, and expertise. According to Gartner, CEOs will be held personally accountable for cyber incidents in the future (Moore 2020).

Awareness to defeat social engineering: All kinds of phishing attacks or attempts of illegitimate access to systems where the human nature is exploited can best be prevented by means of education and awareness campaigns. The human weakness has received more attention in recent years: it is reported that between 80% and 90% of all cyberattacks start with the "human factor." It is therefore important to make employees the most effective firewall of an organization. An important insight is to make employees concerned on a personal level about the devastating impact of a small lapse or carelessness.

Blue and red team exercises: To increase the awareness and to actively take on the perspective of an attacker, role plays with two competing teams—red for attack and blue for defense—have proven to be an effective measure. The incentivization for the players is high and the outcome is an improved secure posture.

Network segmentation and security-by-design: System architectures and networks shall be structured in a way that a breach in one component shall affect as few other components as possible. An attacker shall be prevented from moving laterally in a system or network. Firewalls and explicit transitions between network segments and avoiding reuse of credentials across applications and devices constitute examples for a basic cyber hygiene.

Security-by-default: The configuration of all critical components in a network or organization shall be such that no further adaptation is necessary in order to guarantee secure operation. For instance, devices shall only be connectable to a network or ecosystem, *after* an individual strong password is pre-set or chosen by the owner.

5 Why the Quest for Efficiency Proves to Be the Greatest Opponent of Resilience

In contrast to resilience, which we defined above as being adaptable to disruptive changes in the environment, efficiency can be defined as being adaptable to an existing environment. More precisely, in many situations, efficiency translates to a lean way of working in a given environment, in which the essence of a certain goal becomes the focus of our actions. Every detour, no matter how small, should be avoided in favor of an optimized cost-benefit ratio. More precisely, those detours usually provide no benefit, but instead come with a cost. The primary goal when striving for efficiency therefore often consists of reducing cost, not increasing value or benefit.

Reducing cost comes in a number of facets: it means reducing the overhead of managing more than what is utterly necessary and abandoning buffers and additional resources, in particular those that seem superfluous in the here and now. It therefore also means decreasing diversity and redundancy, the two key features of resilience.

From a short-term perspective, this train of thought seems understandable; reducing slack easily finds economic justification in the medium term. However, it excludes from our view all scenarios that are referred to as black swans or the unknown unknowns. These scenarios often come with risks that are intractable in practice. The likelihood is low, but the impact could be devastating. Out of an obsession with such short-term efficiency and growth, humans tend to become blind to such severe scenarios that come with low probability. M. A. Goldberg in his 1975 article "On the Inefficiency of Being Efficient" calls it the narrowness with which problems are defined and the equally narrow range of alternatives sought for solution. Efficiency strives for short-term optimization, whereas resilience seeks long-term optimization.

What if the relentless pursuit of efficiency, which has dominated American business thinking for decades, has made the global economic system more vulnerable to shocks? This thought was brought up by Prof. William Galston in a March 2020 *Wall Street Journal* article on the Corona pandemic. The Corona crisis, which, if understood as a fairly unexpected and unlikely event with devastating consequences, should force us to rethink our security posture. Lean manufacturing and minimal "just-in-time" inventories have demonstrated how Corona has left many organizations unprepared, be it for the weekly shopping or the supply in hospitals.

When efficiency is understood as concentrating all our activities to a single geographic region or to using a single specific tool, all redundancy and slack is designed away. With such single point of failure, the insurance of diversity and redundancy is lost. An extreme local weather phenomenon or a tenacious virus could bring the global economic value creation to a halt. With increased efficiency, we end up with declined resilience.

Martin (2019) gave an illustrative example of declined resilience in monocultures: almonds were once grown in a number of places in America. Today, more than 80% of the world's almonds are produced in California's Central Valley. The high efficiency is a high price at the same time: California's almond blossoms all need to be pollinated in the same narrow window of time, because the trees grow in the same soil and experience the same weather.

This circumstance is exemplary for monocultures: a single factory produces a product, a single company holds sway in an industry, a single piece of software dominates all systems, a single sensor is responsible for two horrific aircraft crashes and years of grounding, and a single state has de facto global sovereignty over certain IT products and services. Diversity declines. Redundancy declines. Resilience declines.

Yet shouldn't resilience be a public good, in particular in the face of unexpected shocks? Shouldn't our society have the ability to recover from difficulties, to spring back into shape after an unexpected shock? Isn't resilience a fundamental societal need? Surprisingly, the seemingly most resilient form of governance, for which also Churchill has not found any better alternative, is highly inefficient: democracy. There are other examples indicating that thousands of years old achievements and circumstances prove either resilient or efficient, but not both. As Livnat and Papadimitriou have illustrated in 2016, sexual reproduction favors genes that work well with a greater diversity of other genes, which makes the species more adaptable to disruptive environmental changes—exactly what we defined as resilience. The efficiency of the selection process, however, is remarkably low when it comes to approximating the optimum solution. Nature prefers long-term over short-term optimization. *Tinder* prefers short-term efficiency over long-term resilience. Its efficiency lowers expectations by both parties, which is why around 32% of the couples who met online had broken up after a year (compared to 23% of those who met offline), according to a 2014 study by Michigan State University.

6 Resilience in Cyberspace, But How?

How then do we achieve more resilience, possibly without having to sacrifice the efficiency we have come to love, the efficiency that serves as a key source of competitive advantage?

Finding the sweet spot between two competing properties, resilience and efficiency, turns out to become a key challenge in the future of designing robust IT systems. For example, in data-driven economies, sorting algorithms play an important role in handling huge amounts of data. These algorithms are often selected solely based on their performance. As David Ackley reported in 2013, the most efficient algorithms known today, *MergeSort* and *Quicksort*, are not resilient against errors in the underlying comparison function. The slightly less efficient *BubbleSort*, however, is robust thanks to its redundancy in the comparison function. This example shows a viable path by accepting minimal losses in efficiency but, at the same time, increasing resilience in return.

Roger L. Martin proposes a number of steps to foster resilience. For example, the increase in efficiency as a justification for mergers in antitrust policy has led to the concentration of powers and market domination by a few powerful players. Such concentration leads to unstable monocultures and should therefore be avoided—in particular if local scopes and policies are dominated by global ones. The philosophy behind entrepreneurial decisions should not be guided not only by the quest for efficiency only but also by considerations of resilience.

Another example is the deliberate introduction of friction. The injection of productive friction can serve as a regular immunotherapy to build up resilience of organizations, societies, and ecosystems. Friction comes in various flavors, for instance as early as in the careful selection of candidates during the hiring process. Investing in a careful recruitment journey proves key to creating sustainable work environments: preserving and developing an organization's culture not only creates resilience but also long-term productivity. In general, labor should not be considered a cost; it is a productive resource that should be treated as such. For example, allowing for slack so that employees have time to serve customers in unanticipated, yet valuable, ways creates enthusiasm and excitement, both for customers and employees.

Coming back to connected ecosystems and distributed networks, a number of principles can help increase resilience. Most important is a mind shift of participants and operators to move away from striving for efficiency. Instead, parts of the resources saved through digitization should be spent on increasing their resilience. Slack and friction should be introduced at every node of the network, well accepting that the increase of the cost per node is unavoidable to preserve the value of a network as described by Metcalfe, Law, and others.

In terms of information security, every node's effort boils down to having strong authentication processes (with more than a single factor) and a zero-trust mindset (with more local checks and skepticism than stupid belief). On the network level, it is important to implement compartmentalization (with more containers than one—possibly efficient—monolith) and redundancy (with more than the absolute minimum necessary). Finally, we should remind ourselves that richness in species (or "biodiversity") is a desirable goal also in the digital space.

Success Factors for Cybersecurity

- Network effects are responsible for 70% of the value created by technology, from which also attacking adversaries benefit. Including the measures necessary for protecting a network's value, the cost of a network grows beyond linear.
- A host of classical cybersecurity measures exist for keeping attackers at bay, e.g., continuous monitoring and anomaly detection, incident response and digital forensics, regular penetration testing, creating awareness, blue and red team exercises, network segmentation and security-by-design, and security-by-default.
- The ubiquitous quest for efficiency is a key reason of why cyber insecurity is on the rise. Consequently, efficiency as a business goal should be rethought and put in comparison with a potential increase in resilience for socio-economical ecosystems.

References

Allyn B (2020) Researchers: nearly half of accounts tweeting about coronavirus are likely bots. https://www.npr.org/sections/coronavirus-live-updates/2020/05/20/859814085/researchers-nearly-half-of-accounts-tweeting-about-coronavirus-are-likely-bots. Accessed 5 Sept 2020

Chaffee I (2016) Real or not? USC study finds many political tweets come from fake accounts. University of Southern California. https://news.usc.edu/110565/rigged-usc-study-finds-many-political-tweets-come-from-fake-accounts/. Accessed 3 Sept 2020

Martin RL (2019) Competitive strategy: the high price of efficiency. https://hbr.org/2019/01/the-high-price-of-efficiency. Accessed 8 Nov 2020

Moore S (2020) Gartner predicts 75% of CEOs will be personally liable for cyber-physical security incidents by 2024. https://www.gartner.com/en/newsroom/press-releases/2020-09-01-gartner-predicts-75%2D%2Dof-ceos-will-be-personally-liabl. Accessed 10 Sept 2020

Nakamo S (2009) Bitcoin: a peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf

Newberg M (2017) As many as 48 million Twitter accounts aren't people, says study. https://www.cnbc.com/2017/03/10/nearly-48-million-twitter-accounts-could-be-bots-says-study.html. Accessed 5 Sept 2020

NfX (2017) 70 percent of value in tech is driven by network effects. https://www.nfx.com/post/70-percent-value-network-effects/. Accessed 5 Nov 2020

Varol O, Ferrara E, Davis CA, Menczer F, Flammini A (2017) Online human-bot interactions: detection, estimation, and characterization. https://arxiv.org/pdf/1703.03107.pdf

Young VA (2020) Nearly half of the twitter accounts discussing 'reopening America' may be bots. https://www.cmu.edu/news/stories/archives/2020/may/twitter-bot-campaign.html. Accessed 4 Nov 2020

Raphael M. Reischuk, Dr., is distinguished consultant for information security, head of cyber security, and Partner at Zühlke, Vice President of the Cyber Security Committee of Digital Switzerland, and a member of the Cyber Security Advisory Board of the Swiss Academy of Engineering Sciences (SATW). Raphael appears regularly on topics of network, web, and cyber security. He is the author of numerous scientific publications in various fields of information security and cryptography, for which he has received several awards. After studying computer science and mathematics with a focus on information security, Raphael Reischuk received his PhD in web and cloud security at the CISPA Helmholtz Center for Information Security, Germany, and Cornell University, NY, and graduated with distinction. At ETH Zurich, he has done research and teaching on secure Internet protocols and co-developed the new Internet architecture SCION. After a comprehensive research work spanning 12 years, he currently applies his expert knowledge as a distinguished consultant at Zühlke.