The interplay between decentralization and privacy: the case of blockchain technologies

Primavera De Filippi
CNRS — Berkman Center for Internet & Society at Harvard

Introduction

With the current state of telecommunication technologies, it is becoming harder to communicate on the Internet without leaving traces or disclosing information to centralized third parties –be they either governmental agencies or corporations (Lyon, 2014). Indeed, in spite of its distributed infrastructure, today's Internet is highly centralized. A large majority of the Internet traffic is routed through a few centralized services, controlled and governed by a few large corporations. Centralized platforms are useful coordination tools, which provide end-users with great comfort and convenience. Yet, they often come at the expense of privacy and individual autonomy, as users delegate to third parties the task of managing their online activities (De Filippi, 2013). Moreover, most centralized platforms rely on unifying network points that can be regarded as single points of failure, to the extent that they are more likely to be attacked by malicious users, or simply be coerced by governmental agencies in order to disclose information about specific users (Schneier, 2009).

There is a growing interest in decentralized architectures as a way to protect one's privacy against the growing authority and surveillance of centralized third parties. Decentralized architectures are much more supportive of individual freedoms, such as privacy and freedom of expression (Ziccardi, 2012). Yet, although less likely to be subject to centralized control, they are much more difficult to implement as effective coordination mechanisms. In particular, decentralized architectures suffer from an important drawback that might ultimately impinge upon users' privacy: if the price of centralization is trust —since one needs to trust the centralized party to act in compliance with privacy rights— decentralization comes at the costs of transparency —as coordination amongst a distributed network of peers is generally achieved through the disclosure of everyone's interactions.

While decentralized architectures can provide more privacy at the content layer (to the extent that content is properly encrypted), they cannot protect themselves against third parties analysis of metadata, which is openly disclosed to every node connected to the network. Unless additional technological means are employed to protect the privacy of metadata, it might therefore turn out that highly decentralized infrastructures, designed to promote privacy and autonomy, end up

being more vulnerable to governmental or corporate scrutiny than their centralized counterparts.

This paper will focus specifically on blockchain technologies, as an example of extreme decentralization which “suffers” from radical transparency, and whose privacy-enhancing features might therefore be turned against a decentralized network of users (Bradbury, 2013). Our argument will focus on the novel capacities of data mining techniques, as applied to the analysis blockchain metadata, and how these could be just as invasive as traditional surveillance on centralized platforms. We will then move on to analyse how recent advances in cryptography might potentially resolve the inherent trade-off between privacy and transparency, for the ultimate benefit of end-users.

I. Centralized online architectures

The Internet was originally designed as a highly decentralized infrastructure —a network run by everyone and owned by noone. Over time, as new commercial opportunities for Internet services emerged, the Internet grew into more and more centralized clusters (or walled gardens) governed by a few large corporations, *e.g.* Google or Facebook. These services operate in a distributed manner (*i.e.* they use decentralized infrastructures), yet, their governance model is highly centralized. Despite the benefits it provides in terms of coordination and control, the concentration of power in the hands of a few online service providers is progressively leading to a situation of ubiquitous surveillance.

A. Better control and coordination

Coordination can be easily achieved in centralized systems, where information is routed through a series of trusted nodes that collect information needed to coordinate network’s activities. Information is processed centrally and then dispatched to each individual user on a selective basis, *i.e.* only inasmuch as necessary to ensure the proper operations of the network. Centralized coordination thus provide two important benefits: first, it reduces the number of transactions (and the transaction costs) necessary to coordinate a disparate group of individuals; secondly, it reduces the amount of unnecessary disclosure that users would otherwise have to cope with, in a more decentralized system.

The drawback is, of course, that centralized coordination comes at the cost of entrusting a centralised third party with managing all users’ activities and communications (Duffany, 2012). Indeed, to the extent that these trusted third parties are in charge of both holding and managing information on the behalf of all users in the network, each user must trust them to both delegate tasks effectively and only use their personal information for the appropriate ends. In short, greater centralization is often accompanied by a greater need for users to blindly trust a central authority (Bilder, 2006).

Yet, sometimes, these centralized authorities are not worthy of trust. The reason is that centralized coordination could potentially lead to some abuses of power by large online operators, whose (economic) interests are often improperly aligned with that of their user-base. However, to the extent that they control the operations of the network, online operators are often tempted to impose a certain number of restrictions (or obligations) on all users participating to the network, which are forced to accept these conditions in order to benefit from the service, even if this goes against their own interests. This is, for instance, the case of most large online operators, such as Google, Twitter, or Facebook, which can leverage their bargaining power in order get users to accept specific Terms of Use which are much more favorable to the online operator than the users (Chiu, 2011).

Beyond the privacy concerns that this entails (described in more details below), centralization also provides more room for top-down regulation and control. Regulation within a centralized system is facilitated by the fact that it is easy for centralized authorities to monitor everything that is going on in the network. As a result of their privileged position, the operators of a centralized platform can intervene immediately, if there appears to be something wrong. Yet, to the extent that they have full control over the operations of the platform, these operators could also unilaterally decide to reprimand, punish or ban ill-intentioned individuals, insofar as they do not behave according the rules. Perhaps even more critical is the fact that online operators increasingly rely on technological means in order not only to monitor, but also dictate how people can or cannot interact with their platform — thereby significantly reducing the autonomy and impinging upon the privacy of users.

B. Implications on privacy and autonomy

As technology develops, new opportunities are offered to us in terms of communication and information sharing. Yet, to date, our interactions are for the most part mediated by a variety of connected devices that communicate information to one or more Internet service providers.

To the extent that they collect relevant data concerning users' activities and online communications, centralized online platforms constitute very valuable sources of information, which can be exploited by both ill-intentioned individual (*i.e.* hackers) and governmental agencies. Besides, most centralized operators are subject to the regime of intermediaries liability limitations, designed to promote cooperation between the government and online operators (Peguera, 2009) by encouraging them to disclose information about alleged infringers in order to escape from potential liability claims.

Although it is always possible to cloak ourselves by means of specific technological tools designed to obfuscate the content or the source of online communications, most Internet users have nowadays lost not only the technical ability, but also the willingness to keep their personal data privately and securely stored in a trusted datastore. In particular, in many cases, surveillance has become a precondition for

users to enjoy a more personalised service. Online operators are thus legitimized for monitoring the habits and learning about the preferences of their user-base, asking users to disclose more and more personal information in order to better profile them and —ideally— reward them with a service that is (allegedly) more fine-tuned to their respective preferences and needs.

Thus, in spite of the privacy-enhancing technologies and applications deployed in order to help users protect personal data and confidential information against the constant probing of established powers (e.g. including anonymous proxies, VPNs and encryption mechanisms), none of these applications actually succeeded in counteracting the more general trend towards surveillance and control.

Indeed, as both automation and centralization increase, technological advances make it easier for institutions to control the flow of information, monitoring our everyday activities and keeping track of everything we do online and offline (Lyon, 2001, 2002, 2003, 2004, 2014).

This situation of radical transparency is, however, characterized by strong asymmetries of power (Allmer, 2012; Fuchs, 2012). While users cannot protect their personal data against the eyes of centralized online service providers and governmental bodies, the reverse is not true.

Large online operators, such as Google, Facebook, Apple and many others constantly collect data provided —either willingly or unwillingly— by their user-base. Such data is subsequently aggregated, analysed, interpreted or otherwise processed with a view to provide users with a more customized service (Dwyer, 2011). But the algorithms subtending the processing of such data are never disclosed to the public. They are generally kept secret, for the purpose of maintaining a competitive advantage over other intermediaries (Latzer & al., 2014), but also under the assumption that, if they were publicly disclosed, it would be too easy for people to cheat or simply bypass them. As a result, users have no way of identifying the manner in which their data is being processed, and therefore to asses the impact of such processing over their own life (Sandvig, 2013; Bozdag, 2013).

In addition to the implications on users' fundamental right to privacy, the *modus operandi* of these centralized intermediaries is also likely to hinder the autonomy and the individual agency of end-users. As Frank Pasquale has eloquently stressed out, we are now living in a 'black box society' (Pasquale, 2015), where powerful interests increasingly rely on secrecy not only in order to increase their profits, but also as a means to control the way in which individuals can (or cannot) act.

II. Decentralized online architectures

As a reaction to the growing centralization of data in the hands of a few large online service providers, decentralized and federated initiatives have emerged in recent years (Aberer & Hauswirth, 2002). Examples of federated platforms include the distributed social network Diaspora, the collaborative editing platform Kune, and the newly released Federated Wiki. Decentralized networks include the peer-to-peer file

sharing networks BitTorrent, the decentralized communication system FireChat, and the decentralized payment system Bitcoin, based on blockchain technologies.
In addition to being more censorship-resistant, these platforms are also more likely to protect the privacy and confidentiality of information since there is no centralized intermediary that controls all the information flows (Agre, 2003). Yet, given that no central entity is in charge of coordinating users behavior, information needs to be disclosed to a distributed network of peers to effectively align actions in the network (Galloway, 2004). Decentralized networks thus require more transparency in order to effectively coordinate activities between the network nodes.

A. Impact on privacy and autonomy

The privacy of communications can be jeopardized in a variety of ways, depending on the types of architectures at hand. In most centralized systems, users do not need to worry about securing their own communication channels, which are usually managed by a central (trusted) authority. Yet, given that this central authority has complete access to everyone's communications, surveillance remains, almost inevitably, an most important threat to privacy. In more decentralized systems, surveillance is more difficult to achieve (although not impossible) because there no single entity that controls and manages everyone's communications. However, to the extent that users need to secure their own communication channels, bad securitization practices and security flaws in client-side software become much more relevant, especially when data is stored on users' devices (Cole, 2011).
This notwithstanding, most of the decentralized architectures that we encounter today are actually intended to preserve and to promote user's privacy by focusing on at least one of two different privacy paradigms: *confidentiality* and *control*. The former is meant to ensure the confidentiality of people's personal data, by ensuring that their interactions and online communications are shielded from the eyes of third parties. People are thus granted a means to escape from the surveillance of governmental agencies and from the ubiquitous data collection of existing commercial offerings. This is the case, for instance, of tools like TOR and FireChat, which are used by several dissidents around the globe, in order to self-organise outside the purview of the state. The latter is designed to enable people to more actively decide when and with whom to share their own personal information. Rather than focusing on the concealment of personal data, it empowers individuals with a greater degree of control over the collection and use of their data in the context of their ongoing interactions with third parties online operators. Initiatives of this kind include, for instance, the various personal data store initiatives, such as *Personal Black Box* and ID3's *Open Mustard Seed*.
When it comes to autonomy, decentralized architectures also have an important role to play, by eliminating user's dependency towards centralized online operators. This is well illustrated by the FreedomBox initiative, a "personal server" aiming to preserve privacy and individual autonomy by providing a secure platform for personal data storage and applications deployment.

B. The challenges of decentralized coordination

Beyond the advantages it provides in terms of privacy and autonomy, decentralization also raises a few challenges when it comes to the coordination of a large number of individuals. In centralized systems, coordination is easy as there exists one central authority in charge of coordinating the behaviors and the activities of every participant in the system. The ease of coordination has, however, to be counterbalance with the fact that centralized coordination necessarily, and inevitably requires people to trust a centralized third party to behave in their own interests. In the context of decentralized systems, coordination needs to be done in a more distributed manner, as there is no central entity that is aware of what everyone else is doing. Moreover, as a result of the greater autonomy provided to every node in the system, it is difficult to foresee and regulate the behavior of these nodes, which are free to act in a way that does not require previous approval by anyone.

Federated architectures, such as Diaspora, Kune, or the Federated Wiki, constitute a hybrid solution between centralized and purely decentralized architectures. Just like centralized systems, they requires user to trust specific entities (or hubs) which are in charge of managing their own network. However, given that is more than one hub in the system, users are free to interact with the one they feel the most comfortable with, or even to create their own hub. Each hub manages its own internal network through a centralized coordination system, whereas all external coordination is done through a system of peer-to-peer communication among hubs —which only need to disclose the information that is strictly necessary to ensure the operations of the network as a whole.

At the end of the spectrum towards decentralization are pure P2P networks like Freenet, Tor or BitTorrent. These fully decentralized systems make coordinating behaviour within the network much more difficult. Without any centralized third party to rely upon, decentralized platforms need to find new ways to effectively coordinate behavior among a disparate network of peers, without passing through any trusted authority (Oram, 2001). This is generally achieved by disclosing specific data (or metadata) to every node of the network, so that they can coordinate actions between themselves in a distributed and decentralized manner (Aberer & Hauswirth, 2002).

Specifically, in a centralized platform, communications are first communicated to a centralized —and ideally trusted— authority, which is responsible for dispatching the information to all relevant nodes with the necessary credentials. Conversely, in a decentralized peer-to-peer platform, given that there is no trusted party to rely upon, every node needs to communicate with (and therefore also trust, even if to a lesser extent) every other node in the network, in order to ensure that proper coordination is achieved. Most importantly, given that ill-intentioned users might be

tempted to ‘cheat’ the system, in absence of a central authority in charge of policing the network, there needs to be a mechanism for the network to collectively verify the legitimacy of every individual transaction —which require a high level of transparency in the network. This need for “radical transparency” is well illustrated by the case of blockchain technologies such as Bitcoin for instance, where everyone gets to see what everyone else is doing, and it is up to every node in the network to verify that everyone else is acting in accordance with the rules of the system (Bradbury,2013; Nakamoto, 2008).

Thus, it appears that, the more we shift towards a decentralized infrastructure, the less we need to rely on trust and the more we rely on transparency instead. If decentralization can contribute to promoting user’s privacy and autonomy, it might, however, come at the cost of radical transparency.

III. The Interplay between Transparency and Privacy

If decentralized peer-to-peer technologies do, indeed, require significant transparency in order for users to interact in a coordinated way, decentralization might have both positive and negative implications for the protection and preservation of users’ privacy. On the one hand, it might improve the ability for users to control the use and dissemination of their personal data, by making them less dependent on centralized third parties or service providers that own and control most of the data provided by their user-base. On the other hand, the degree of transparency required for the purpose of coordinating the activities and communications of a large network of peers might reduce their ability to protect their own data —or metadata— from the eyes of (potentially malicious) third parties. Hence, when analysing the privacy implications of decentralized architectures, we must understand whether the privacy gains resulting from decentralized coordination in a distributed network are greater than the costs derived from the mandatory disclosure of metadata it entails.

A. Transparency at the detriment of privacy?

As a general rule, decentralized platforms tend to be more respectful to users’ fundamental right to privacy to the extent that they make it harder for centralized third parties to have complete control over users’ data. In a peer-to-peer system, data is stored either locally on the users’ devices (and communicated only to the relevant parties) or is spread all over the network where multiple parties (or peers) are in charge of storing and processing small fragments of such data. While the information is theoretically visible to anyone, the use of end-to-end encryption allows users to communicate privately between one another, without having to entrust anyone with the task of managing and transferring personal information (Clarke & al., 2002; Cutillo & al., 2009).

However, since there is no centralized entity in charge of collecting, storing and processing data, the only way for users to coordinate themselves in a decentralized manner is to make data available to every other user of the network. While the content of communications can be encrypted so that it can only be accessed by the persons to whom it was actually addressed, the metadata related to these communications *(i.e.* who is talking to whom, for how long, and what is the type of transaction in which they participate) needs to be visible to all (Abiteboul, & Marinoiu, 2007).

Of course, many exceptions to that rule exist. The *Onion Router* (TOR) is a clear example of how a decentralized network can be designed so as to preserve anonymity with regard to the source of online communications. Similarly —although less widely adopted— the *Freenet* project implements a distributed anonymous information storage and retrieval system, whose communications are encrypted by default, and metadata information is obfuscated via sophisticated routing techniques.

Yet, building decentralized systems is, as a general rule, much harder than building a centralized platform. Most of the time, the design challenges of building a decentralized architecture are made easier by giving up the privacy of metadata. While this is not an inherent requirement of any decentralized system, it is, in practice, the most common implementation thereof.

Accordingly, to the extent that decentralized networks require a greater level of transparency than their centralization counterparts, there might be an impending tension between degree of decentralization that a network enjoys and the amount of privacy that users of that network might effectively enjoy (Filipovikj & Holmstedt, 2013).

B. Reduced information asymmetries

Independently of the issues it might raise in terms of privacy and confidentiality, radical transparency could also contribute to reducing the degree of information asymmetries that characterise the relationship between online service providers and their user-base.

There is, indeed, a fundamental difference between disclosing information, subdivided into small chunks, to a distributed network of peers and disclosing it, in whole, to a centralized entity. The former requires users to entrust a large number of individuals with only small chunks of personal data, which cannot be easily put together without incurring extensive coordination costs. The latter presupposes that one centralized entity be intrusted with all data stemming from every node in the network and —insofar as they rest into a central repository— these data can ultimately be aggregated into a larger dataset, cross-referencing the data of multiple users, at virtually no costs. Hence, decentralization might significantly reduce the information asymmetries (and the resulting power asymmetries) that are generally enjoyed by most centralized service providers (Themelis, 2013).

Perhaps most importantly, transparency might have a considerable effect on the asymmetric power dynamics associated with the 'black-box society' that is slowly emerging on the Internet (Pasquale, 2015). As mentioned earlier, most of the centralized online platforms, such as Google or Facebook, increasingly rely on algorithms whose rules are often not publicly disclosed and —even if they were— would probably be too complicated for many individuals to understand fully. The complexity and the opacity of these algorithmical rules make it very difficult for users to understand they way in which their life is actually being affected by these online platform. Besides, by not disclosing the algorithms that underpin the services they provide, these online operators preserve the capacity to unilaterally modify their offer at will, without the approval of their user-base.

A recent example of the tremendous power wielded by centralized intermediaries is Facebook's manipulation of NewsFeed in order to gauge emotional reactions of users. As part of a social experiment— Facebook arbitrarily modified the NewsFeed displayed to a certain categories of users, in order to identify whether (and how) their mood would be affected by different types of content. The experiment has caused strong controversy on the moral and ethical aspects of such undertaking (Puschmann & Bozdag, 2014). Mostly importantly, it has shown that Facebook has thus far acquired enough power to effectively control the mood, and perhaps even the actions of individual users, without giving users the possibility to even know what is happening behind the scenes .

In a decentralized network, a similar operation could not be achieved without first obtaining the consensus of the network as a whole, or a very large part thereof. Indeed, given the transparency that is inherent to these network, anyone can monitor the traffic of information in order to make sure that information flows properly to the people it was intended for, and that such information is not being used in a way that might go counter the interests of end-users (Hugues & al., 2006). This ties in with what Hellen Nissenbaum has defined as "contextual integrity" —*i.e.* the right not to control personal information, nor to have access to restricted or confidential information, but rather the right to live in a world in which our expectations of privacy (with regard to the flow of personal information) are, for the most part, met. Expectations, in this sense, do not refer only to the customary understanding of how is personal data currently collected or processed by online operators, but rather to the way it *should* be treated, according to social, moral and political norms.

C. Modified power dynamics

Decentralization brings along a shift in the power dynamics that characterise the traditional relationship between centralized online operators and their users. Insofar as the service operator controls the platform on which users interact, it has the ability to control and regulate most of their online activities —whether this is not by contractual means (e.g. end-user licensing agreements) or technical means (i.e. as a result of technical design). In a decentralized context, users activities are

governed exclusively through the code or the protocol of the underlying network or technology, and nothing else. Yet, as opposed to more centralized systems, no one can impose, or even modify any of these technical rules, since the software must be run directly from the users' computers or devices. It might seem, therefore, that decentralization is likely to promote individual freedoms and emancipation. However, given the lack of supervision or oversight, decentralized technologies are also more likely to be co-opted by established powers in order to further their own interests.

Let's look at the bright side first. As we have discussed earlier, decentralized technologies offer important advantages to end-users. Not only do they provide for greater privacy and confidentiality, but they also promise a greater degree of autonomy to end-users, who are no longer subject to the terms and conditions unilaterally imposed to them by large online operators. While centralized platforms are generally operated by commercial players whose interests are often incompatible with that of end-users, decentralized platforms are operated directly *by* the community *for* the community; and are therefore designed in such a way as to best fulfill the community's needs (Oram, 2001).

Decentralized networks are also generally more open. Given that there is no centralized authority to verify user's credentials, anyone is free to join the network at any time, provided that they comply with the network's protocols and technological standards. The result is a greater degree of interoperability and diversification (*e.g.* multiple clients can implement differents set of features on top of the same protocol layer). And given that user's data is stored on users' devices, decentralization also considerably reduces the risk of user lock-in, which is often due to the lack of data portability.

But, decentralization comes with some downsides. Beyond the issues surrounding coordination and control —which we have already identified earlier— there is also the question of *cooptation* (or perhaps simply manipulation). Given that there is no one in charge of ensuring the proper operation of the network, vested interests might be tempted to jump in and take over the network, or simply influence it —either directly or indirectly— so as to further their own interests. at the expense of the other community members. TOR, for instance, has allegedly been infiltrated by the US government, at several occasions. Infiltration, in this case, simply entails deploying a large number of relays into the network. Indeed, anyone who controls a sufficient portion of the TOR network can get a fairly good overview of the traffic that is being routed through the whole network —and this information can subsequently be used in order to de-anonymize some of the users. Perhaps a better example of evolving power dynamics can be observed in the context of the Bitcoin network. As a decentralized payment system, the Bitcoin network relies on the work of peers connected to the network (so-called *miners*) in order to verify the validity and integrity of every transaction. Given the economic incentives that emerged with the recent rise in value of Bitcoin, and given the open nature of the Bitcoin network, the *mining* of Bitcoin has grown more and more centralized, and the network is now for the most part controlled by only a few large mining pools, which —if they were

to collude— could easily take over the whole network with a 51% attack (Kroll & al., 2013).

The case of blockchain technologies

This paper will focus, in particular, on the technology underlying the Bitcoin network —the blockchain— as a particular example of a techno-political architecture that might considerably affect the privacy and individual autonomy of end-users. Indeed, blockchain technology represents an important and promising development in Internet technologies insofar as it makes it possible for people to transact and to interact with one another without relying on any centralized intermediary. It constitutes, as such, a new means of coordination enabling new forms of collective action amongst disparate groups of peers that were previously thought to be impossible (Swan, 2015).

A. Technical overview

A blockchain is a decentralized ledger (or state machine) that relies on cryptographic algorithms and economic incentives in order to ensure the integrity and legitimacy of every transaction (or state change). A copy of the blockchain is shared amongst all nodes connected to the network, which comprises the history of all valid transactions. Each transaction is recorded into a 'block' which is appended *sequentially* to the previous block of transactions (Nakamoto, 2008). In order to prevent anyone from tampering with past transactions, the blockchain acts as an append-only ledger —*i.e.* once information has been recorded onto the blockchain, it can no longer be edited or deleted. The result is a long chain of blocks (or *blockchain*) that represents the whole chain of transaction ever since the first *genesis block* (Sprankel, 2013).

The blockchain can thus be regarded as a secure database that comprises a public log of all transactions which have been thus far validated by the network. In view of its decentralized nature, the security of the blockchain and the validity of every transaction can only be ensured through distributed consensus (*i.e.* through nodes verifying the integrity and legitimacy of each block, independently of any trusted third party). This requires that the transaction history be made available to the public, so that it can be easily verified by anyone. The consequence is, however, that anyone who has access to a copy of the blockchain also has access to the current (and past) consensus state —with regard to *e.g.* the flow and the amount of all validated transactions (Ober & al., 2013).

The main innovation beyond the blockchain is the ability to validate transactions in a decentralized manner, without relying on any trusted authority. Until recently, a decentralized currency that operates independently of any financial institution was simply inconceivable. Blockchain technologies make it possible by allowing for transactions to be verified and computer logic to be executed in a decentralized

manner. Instead of requesting confirmation for every transaction to a centralized authority, the blockchain's distributed consensus is such that any attempt at tampering with the consensus state will most likely be rejected by the network as an invalid transaction.

Initially developed as part of the Bitcoin network, the blockchain is a general purpose technology that can be used for many other kinds of applications which formerly required the existence of a trusted third party: from decentralized domain name systems (*Namecoin*) to decentralized land and commercial registries (*Factom*) or any decentralized application that can be run on the *Ethereum* blockchain. Many things that previously required a centralized intermediary to coordinate the action of multiple people can now be done in a decentralized manner via the blockchain. As such, the blockchain constitutes yet one more step towards the process of disintermediation.

B. Privacy, anonymity and pseudonymity

It is often believed that —because of their decentralized nature— blockchain technologies might contribute to promoting individual's privacy and autonomy. Indeed, instead of relying on the coordination activities of a centralized authority, the blockchain operates through a decentralized public ledger which is regulated exclusively by code and algorithmical rules. Yet, in order to allow for meaningful coordination, the blockchain must be both accessible and auditable by every node in the network. Indeed, without a centralized intermediary, the only way for individuals to properly coordinate themselves is for everyone to share a common datastore with the most updated state of the consensus. The inherent transparency of blockchain technologies is thus a necessary condition to successfully coordinate the behavior of several individuals that do not know (nor trust) each others.

Such degree of transparency might not always be desirable. In many cases, the transparency inherent to these technologies actually goes counter to the traditional expectations of privacy. Yet, the fact that the blockchain cannot have privacy does not necessarily means that also its users cannot.

Bitcoin and many other blockchain-based applications mitigate the costs of transparency by virtue of anonymity and/or pseudonymity. For example, Bitcoin is often described as an anonymous decentralized cryptocurrency, in that it allows people to transact with one another without having to disclose any information related to their actual identity. Since the public addresses used in every transaction are random numbers that do not need to be associated with an identity, even though the transaction history is made publicly available to anyone, people's privacy can be preserved as long as it is not possible to trace back these transactions to any given identity. Yet, the truth is that Bitcoin is actually not anonymous —but rather pseudonymous— and that the inherent transparency of the Bitcoin blockchain might ultimately hinder —instead of furthering— the privacy of end-users.

Indeed, while it might provide a limited degree of privacy, pseudonymity is far from being able to ensure a reliable amount of privacy protection. Specifically, anonymity is such as to make it impossible to relate multiple transactions with a single source or destination, whereas pseudonymity only implies that the identity of the person(s) associated with that specific source or destination cannot be (easily) established. Although blockchain technologies could potentially be implemented in an anonymous way, most of the blockchain-based application implemented so far do not provide strong anonymity support (Reid & Harrigan, 2013).

For instance, the design of the Bitcoin blockchain is such that the more an address is used, the more information can be inferred from this address. While good privacy norms would require people to constantly generate a new address before performing a new transaction, only a minority of people actually engage in these practices. In the Bitcoin space, most non-tech savvy people simply reuse their Bitcoin address without realizing that, by doing so, they are publicly disclosing valuable personal information. This can be quite problematic from a privacy standpoint. Given the transparency and non-repudiability of the Bitcoin blockchain, it is possible to keep track of every transaction involving a particular Bitcoin address. Regardless of how careful a person has been to hide his or her identity in the past, once the identity of the person owning that Bitcoin address has been established, it then becomes possible for anyone to retroactively associate to that person all the transactions which have previously been made to and from that address (Moser, 2013).

In this regard, specific data analysis techniques can and have already been deployed to extract, deduce or infer new information from the transaction history that has been stored on the blockchain.

C. Blockchain analytics

As Bitcoin adoption grows and expands into more and more regulated sectors of activities, the ability to identify the source and destination of financial transactions becomes an even stronger imperative. Indeed, the finance industry is a strongly regulated industry, which needs to comply with significant formalities in order to ensure that it is dealing only with legitimate clients. Given the difficulty to establish the identity of pseudonymous blockchain addresses, the Bitcoin blockchain has often been used for illicit activities —see, *e.g*. the case of the Silk Road market-place (Barratt, 2012). Although these constitute only a marginal portion of all Bitcoin transactions, the risk to be regarded as facilitating criminal activities might be sufficient to dissuade financial or commercial operators from interacting with non-identified Bitcoin addresses (Moser & al., 2013).

In view of the growing commercial impact of Bitcoin, many companies are turning blockchain analytics into a new business model, providing tools for other companies to comply with the law, such the anti-money laundering (AML) regulations. By associating pseudonymous Bitcoin address with real-world entities, these tools

identify the list of Bitcoin addresses which are knowingly related to criminal activities, and which should therefore be blacklisted by any law-abiding operator.
For instance, companies such as *Coinalytics*, *Coinometrics*, etc. are building tools for people in the Bitcoin industry to extract new and meaningful insights from the Bitcoin blockchain, so as to support business intelligence and compliance with the law. On that regard, the company *Elliptic* recently launched a new project —the *Bitcoin Big Bang*— which provides an interactive tool for visualizing past and current transactions on the Bitcoin network, along with the identity of the person or company that issued or received these transactions. Thanks to these services, users can immediately get a good grasp of what is going on in the Bitcoin space, in order to make better informed decisions as to whom they should transact with (Moser, 2013).
All of these initiatives are thus challenging the initial conception of the Bitcoin network as a means for people to bypass traditional financial institutions in order to freely and anonymously transfer value. Indeed, the inherent transparency of the Bitcoin blockchain is such that the history of every transaction can potentially be tracked down, back to the place where it originated (the so-called *coinbase transaction*). This means that any Bitcoin transaction which has ever been issued by an allegedly criminal address, or which has simply transited through a Bitcoin address which is associated to a criminal identity, will be forever 'tainted' by its own history (Moser & al., 2013, 2014).
This has great implications as regards the fungibility of Bitcoin. Given that financial or commercial operators might refuse to deal with these kinds of transactions, two transactions with the same Bitcoin face value might not have the same operational value, depending on whether or not such transactions could be regarded as potentially tainted transactions. And given that , at any moment in time, a particular transaction might eventually become a tainted transaction, *i.e.* whenever a new public address is identified as being associated with criminal activities, it might turn out that younger transactions (with a shorter transaction history) end up being perceived as more valuable —because less risky— than those with a longer transaction history.
Of course, as the number of initiatives concerned with blockchain analytics increases, also the number of cryptographers and developers seeking to elaborate new mechanisms to preserve the privacy and anonymity of blockchain transactions increases.

D. Modern advances in cryptography

In spite of the transparency inherent to the blockchain, specific mechanisms can be deployed to conceal the source and destination of the transactions, as well as the content thereof. While the blockchain does not, as such, provide any kind of privacy protection, it would be a mistake to believe that the transparency required to operate on the blockchain necessarily and unavoidably goes counter to the privacy

of end-users. Radical transparency only subsists at the most basic layer of the blockchain —that which is responsible for applying the distributed consensus algorithm. But nothing prevents us from building additional layers of encryption and/or obfuscation on top of that layer.

In this sense, the blockchain today is not so different from the TCP/IP layer of the Internet network, which relies on a system of public —but not pseudonymous— IP addresses, with regard to both the source and destination of online communications. Initiatives such as TOR and Freenet have specifically addressed this issue, by introducing an additional layer of anonymity on top of the TCP/IP protocol. Similarly, the content of these communications does not have to be neither public nor transparent (*i.e*. clear text), in order for a machine to efficiently route the packets through the network. The development of a public encryption standard (DES) and, in particular, the popularisation of public-key cryptography (RSA) in the mid-90's, have made it possible for people to use the Internet as a public telecommunication infrastructure, while nonetheless being able to preserve the privacy and confidentiality of their communications.

The same can be done with the blockchain. Technologies such as *CoinJoin*, *CoinSwap*, *CoinShuffle*, etc. are designed to mitigate the privacy-drawbacks of the Bitcoin blockchain by means of obfuscation. These mechanisms exploit one of the most basic features of the Bitcoin blockchain (*i.e.* the independent construction of Bitcoin transactions from other transactions) to provide a greater level of anonymity and confidentiality of transactions (Bonneau & al., 2014).

Ring signatures are also gaining popularity in the blockchain space. As a special type of digital signatures, they allow for a group of people to transact with each other, and with third parties, without revealing the link between an individual signature and an individual's public key. More sophisticated systems exist, such as the Zerocash protocol, which extends the Bitcoin protocol with more advanced cryptographic algorithms (based on zero-knowledge proof) in order to enable people to execute direct payments to each other, without disclosing the source, the destination, nor even the actual amount of these transactions. More recently, *Blockstream* has introduced the notion of *confidential transactions* as a means to improve the privacy and security of the Bitcoin network, without introducing any additional cryptographic primitive to the Bitcoin blockchain. Confidential transactions rely on advanced cryptographic techniques (so-called additively homomorphic commitments) in order to provide a means for people to keep the actual amount of their transactions private, while nonetheless allowing for the public network to verify the validity of these transactions (*i.e.* by making sure that the ledger entries add up). When combined with mixing technologies such as *CoinJoin*, these tools could effectively preserve privacy both at the content level (transaction type and amount transferred) and metadata level (source and destination of the transaction).

It is worth noting that these cryptographic techniques —while ensuring that transaction data remains confidential *by default*— are not necessarily incompatible with the notion of transparency. Users retain the ability to uncloak their transaction

data to third parties (such as escrow agents, investors, potential business partners, auditors, authorities and law enforcement).

IV. Conclusion

Centralized and decentralized platform infrastructures have very different implications for the privacy of end-users. In the case of centralized architectures, end-users communicate to each other through a centralized platform, operated by a trusted authority. In order to dispatch the information it receives to all relevant parties, centralized platforms are thus, *by design*, required to collect the information concerning at least the metadata of online communications. In the case of decentralized architectures, end-users communicate directly to one another, without passing through any centralized intermediary. Communications are routed through a decentralized network —*i.e*. one that does not rely on any single trusted authority. The flipside of that, however, is that, in order to transfer the information to the right destination, the metadata concerning every communication needs to be made publicly available to the network as a whole. In this sense, assuming that a centralized platform is operated by a trusted entity, the centralized model is more respectful of people's privacy than the decentralized model. However, if one cannot trust the central operator to fully respect people's privacy, then a decentralized infrastructure might constitute a better choice.

As opposed to centralized systems, which are characterized by strong information asymmetries between the operator and its users, decentralized systems are more egalitarian, to the extent that they bear an equal level of transparency across all the participants in the network. Transparency does not, as such, constitutes a loss of privacy; yet, in order to protect their privacy against the scrutiny of third parties, users need to deploy additional privacy-enhancing technologies (e.g. end-to-end encryption, network obfuscation tools, etc) on top of the platform (Ziccardi, 2012; Milan, 2013).

Of course, complete privacy and anonymity can never be guaranteed. For instance, whenever there is a backdoor or a bug in the technology, even the most sophisticated encryption techniques will be unable to protect users' privacy and identities. More generally, regardless of how much effort has been put into the design of a secure decentralized architecture, there is no guarantee that people's privacy will never be compromised. Indeed, in a decentralized system, the responsibility of keeping data private merely shifts from the operator to the individual user. While the former is more likely to be coerced (*e.g.* by the government) to disclose information about its user-base, the later is more likely to inadvertently disclose or leak information through an improper use of the platform or tools. All in all, any system whose security ultimately relies on encryption technologies can only be as secure as the ability of users to securely manage their secrets (*e.g.* passwords or private keys).

In this article, we focused on the case of blockchain technologies, as a particular example of a socio-political tool that is completely decentralized both at the

infrastructure and governance level. While the blockchain might definitely promote individual autonomy, the transparency of the blockchain also raises important challenges as regards the privacy and confidentiality of transactions. After showing that transparency is a necessary condition to implement a trustless system that does not rely on any central authority or trusted intermediary, we have shown that —while the blockchain requires radical transparency in order to validate transactions— modern cryptographic techniques can be used to prove that a particular transaction is indeed legitimate, without having to disclose the source, the destination, nor the actual content of the transaction. None of the available techniques are mature yet, but it is only a matter of time and engineering to perfect them. Transparency and privacy should, however, not be regarded as being in a fundamental conflict.

References & Bibliography

Aberer, K., & Hauswirth, M. (2002, March). An Overview of Peer-to-Peer Information Systems. In *WDAS* (Vol. 14, pp. 171-188).

Abiteboul, S., & Marinoiu, B. (2007, November). Distributed monitoring of peer to peer systems. In *Proceedings of the 9th annual ACM international workshop on Web information and data management* (pp. 41-48). ACM.

Agre, P. E. (2003). P2p and the promise of internet equality. *Communications of the ACM*, *46*(2), 39-42.

Allmer, T. (2012). Critical internet surveillance studies and economic surveillance. *Internet and Surveillance: The Challenges of Web 2.0 and Social Media*, *16*, 124

Barratt, M. J. (2012). Silk Road: eBay for drugs. *Addiction*, *107*(3), 683-683

Bechmann, A. (2013). Internet profiling: The economy of data intraoperability on Facebook and Google. *MedieKultur. Journal of media and communication research*, *29*(55), 19.

Bilder, G. (2006). In Google we trust?. *Journal of Electronic Publishing*, *9*(1).

Bradbury, D. (2013). The problem with Bitcoin. *Computer Fraud & Security*,*2013*(11), 5-8.

Bonneau, J., Narayanan, A., Miller, A., Clark, J., Kroll, J. A., & Felten, E. W. (2014). Mixcoin: Anonymity for Bitcoin with accountable mixes. In *Financial Cryptography and Data Security* (pp. 486-504). Springer Berlin Heidelberg.

Bozdag, E. (2013). Bias in algorithmic filtering and personalization. *Ethics and information technology*, *15*(3), 209-227.

Chiu, A. T. (2011). Irrationally bound: terms of use licenses and the breakdown of consumer rationality in the market for social network sites. *S. Cal. Interdisc. LJ*, *21*, 167.

Clarke, I., Miller, S. G., Hong, T. W., Sandberg, O., & Wiley, B. (2002). Protecting free expression online with Freenet. *Internet Computing, IEEE*, *6*(1), 40-49.

Cole, E. (2011). *Network security bible* (Vol. 768). John Wiley & Sons.

Cutillo, L. A., Molva, R., & Strufe, T. (2009, February). Privacy preserving social networking through decentralization. In *Wireless On-Demand Network Systems and Services, 2009. WONS 2009. Sixth International Conference on* (pp. 145-152). IEEE.

De Filippi P. (2013), Ubiquitous Computing in the Cloud: User Empowerment vs. User Obsequity, in: Jean-Eric Pelet, Panagiota Papadopoulou (eds.) 'User Behavior in Ubiquitous Online Environments', IGI Global.

Duffany, J. L. (2012). Cloud Computing Security and Privacy. In *10th Latin American and Caribbean Conference for Engineering and Technology* (pp. 1-9)

Dwyer, C. (2011). Privacy in the Age of Google and Facebook. *Technology and Society Magazine, IEEE*, *30*(3), 58-63.
Filipovikj, P., & Holmstedt, C. (2012). Comparison between centralised and decentralised systems and how they cope with different threats. *Abgerufen am*,*22*(08), 2013.
Fuchs, C. (2012). Critique of the political economy of Web 2.0 surveillance. *Internet and Surveillance. The Challenges of Web 2.0 and Social Media*, 31-70.
Galloway, A. R. (2004). *Protocol: How control exists after decentralization*. MIT press.
Hughes, D., Walkerdine, J., Coulson, G., & Gibson, S. (2006). Peer-to-peer: Is deviant behavior the norm on p2p file-sharing networks?. *Distributed Systems Online, IEEE*, *7*(2).
Kroll, J. A., Davey, I. C., & Felten, E. W. (2013, June). The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In *Proceedings of WEIS* (Vol. 2013).
Latzer, M., Hollnbuchner, K., Just, N., & Saurwein, F. (2014). *The economics of algorithmic selection on the Internet*. Working Paper, IPMZ. Zurich,
Lyon, D. (2001). *Surveillance society: Monitoring everyday life*. McGraw-Hill Education (UK).
Lyon, D. (2002). Everyday surveillance: Personal data and social classifications. *Information, Communication & Society*, *5*(2), 242-257.
Lyon, D. (2003). *Surveillance after september 11* (Vol. 11). Polity.
Lyon, D. (2004). Globalizing Surveillance Comparative and Sociological Perspectives. *International Sociology*, *19*(2), 135-149.
Lyon, D. (2014). Surveillance, snowden, and big data: capacities, consequences, critique. *Big Data & Society*, *1*(2), 2053951714541861.
Milan, S. (2013). *Social movements and their technologies: Wiring social change*. Palgrave Macmillan.
Möser, M. (2013). Anonymity of Bitcoin transactions. In *Münster Bitcoin Conference*
Moser, M., Bohme, R., & Breuker, D. (2013, September). An inquiry into money laundering tools in the Bitcoin ecosystem. In *eCrime Researchers Summit (eCRS), 2013* (pp. 1-14). IEEE.
Möser, M., Böhme, R., & Breuker, D. (2014). Towards risk scoring of bitcoin transactions. In *Financial Cryptography and Data Security* (pp. 16-32). Springer Berlin Heidelberg.
Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Consulted*, *1*(2012), 28.
Ober, M., Katzenbeisser, S., & Hamacher, K. (2013). Structure and anonymity of the bitcoin transaction graph. *Future internet*, *5*(2), 237-250.
Oram, A. (2001). *Peer-to-peer: harnessing the benefits of a disruptive technology*. " O'Reilly Media, Inc.".
Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington law review*,*79*(1)
Pasquale, F. (2015). The Black Box Society. *Cambridge, MA: Harvard University Press*, *36*, 32.
Peguera, M. (2009). The DMCA Safe Harbors and Their European Counterparts: A Comparative Analysis of Some Common Problems. *Columbia Journal of Law & the Arts*, *32*, 481.
Puschmann, C., & Bozdag, E. (2014). Staking out the unclear ethical terrain of online social experiments. *Internet Policy Review*, *3*(4).
Reid, F., & Harrigan, M. (2013). *An analysis of anonymity in the bitcoin system*(pp. 197-223). Springer New York.
Sandvig, C. (2013). The internet as an infrastructure. *The Oxford handbook of internet studies*, 86-108.
Schneier, B. (2009). *Schneier on security*. John Wiley & Sons.
Sprankel, S. (2013). *Technical basis of digital currencies*. Working Paper.
Swan, M. (2015). *Blockchain: Blueprint for a New Economy*. " O'Reilly Media, Inc.".
Themelis, A. T. (2013). Information and Intermediation, Abuse of Dominance and Internet 'Neutrality':'Updating'Competition Policy under the Digital Single Market and the Google Investigations (?). *European Journal of Law and Technology*, *4*(3).
Ziccardi, G. (2012). *Resistance, liberation technology and human rights in the digital age* (Vol. 7). Springer Science & Business Media.

Zyskind, G., Nathan, O., & Pentland, A. S. (2015). Decentralizing Privacy: Using Blockchain to Protect Personal Data. In *Security and Privacy Workshops (SPW), 2015 IEEE* (pp. 180-184). IEEE.