Internet Protocol Adoption: Learning from Bitcoin

Position Paper

Rainer Böhme

Westfälische Wilhelms-Universität Münster
Department of Information Systems
Leonardo-Campus 3
48149 Münster, Germany

rainer.boehme@uni-muenster.de

ABSTRACT

This position paper discusses two related questions:

1. Using the framework of network economics, what are success factors behind the adoption of Bitcoin (a cryptographic curreny), and can we copy them for other protocols?

2. Can we design more successful protocols if we have in-band payment mechanisms to internalize the externalities that emerge during adoption and in steady state?

1. INTRODUCTION

The history of the Internet has unleashed unmeasurable engineering effort into the design of new protocols. However, perceived usefulness, as measurable by counting features or gauging elegant design, has not always been a good predictor for adoption. The search for explanations has led to RFC 5218 [10], a comprehensive list of success factors along with case studies of successful protocols. Many of these success factors can as well be framed in the language of network economics [4]. This enables us to seek for solution approaches in this discipline towards the vision of considering the economics of new Internet protocols already at the design stage.

This position paper in particular presents thoughts inspired by the recent success of Bitcoin, a cryptographic currency system. Section 2 recalls the fundamentals by introducing terminology and pointing out selected economic barriers to protocol adoption in general. It also explains Bitcoin to the extent necessary. Section 3 analyses success factors of Bitcoin and comments on their applicability to other protocols. It turns out that important success factors are tied to the nature of Bitcoin as a payment system. This leads to the discussion in Section 4, which concentrates on means of internalizing externalities by protocol design and explores the idea of built-in payment systems. Finally, Section 5 concludes with open problems and items for a research agenda.

2. BACKGROUND

This section is intended to set the stage by recalling basic notions of network economics and by introducing salient features of Bitcoin. Readers familiar with one or the other topic are safe to skip the respective sections.

Accepted to the IAB Workshop on Internet Technology Adoption and Transition (ITAT), Univ. of Cambridge, UK, December 2013.

2.1 Network Economics

Network externalities

In simple terms, a large part of economics studies the actions of autonomous decision makers (called *agents*) who each seek to maximize their own objective function (called *utility*). *Externalities* arise if the action taken by one agent affects the objective function of other agents. *Network externalities* are a special kind of externality. Each agent's action space can be partitioned into actions that involve participation in a specific network and actions that do not. The objective function of any agent i who participates in the network has strictly greater utility if any other agent j participates in the network as well. This justifies statements such as "the value of the network is super-linear in the number of its users". (See [3] for a discussion of the shape of the functional form.)

Adoption decision

At our level of abstraction, participating in a network is equivalent to adopting a (compatible) protocol. Each agent i adopts if his individual benefit of adoption exceeds his individual cost of adoption. RFC 5218 lists this as basic success factor called *positive net value*. However, in the presence of network externalities, the benefit depends on the number of other agents who adopt. Let $q \in [0, 1]$ be the fraction of agents who adopt. If each individual agent's benefit is below the cost as long as q is small, we face a social coordination problem. No agent is willing to adopt alone but all agents could benefit if they collectively agreed to adopt.

Incremental deployability, another success factor in RFC 5218, can be interpreted as a requirement for the benefit to be significant even if q is small. The other basic success factors in RFC 5218, including open code availability, freedom from usage restrictions, open specification availability, and open maintenance processes, are means of reducing the cost of adoption for all agents, thereby facilitating solutions to the social coordination problem.

Business interests may diverge from this objective. Market power permitting, profit-maximizing providers strive to support smaller networks than socially optimal [4]. Artificial barriers to adoption through pricing, obscurity, and property rights are strategic tools towards this end [12].

Timing and uncertainty

Enter time. Costs and benefits influencing the adoption decision do not always materialize at the same point in time.

Chiefly, *switching costs* are one-off costs incurred at the time of adoption, e. g., the time spent to implement and upgrade systems; or, for user-facing systems, the opportunity cost of starting low on the learning curve. Switching costs are often *sunk*, which means that they are not recoverable after a plan change. Likewise, they arise anew when the next protocol adoption is due. Benefits typically materialize over time and may grow in the presence of network externalities as other agents adopt sequentially. Misalignment in the timing of costs and benefits complicates the adoption decision because costs need to amortize over time.

With time comes uncertainty. When operating under uncertainty, agents must anticipate future costs and benefits, typically by treating them as random variables and taking expectations. Since there is always a small chance that a future transfer of wealth will not happen (or will not affect the agent anymore) because the state of the world has changed, it is reasonable to *discount* more distant costs and benefits exponentially over time. Uncertainty adds even more complications if agents are *risk averse*, meaning that they prefer a smaller profit with certainty over a higher expected profit under uncertainty. Even if agents are *risk neutral*, they may prefer to delay an adoption decision in order to *improve their information* about the likely realization of critical random variables. For example, they might wait and see how many other agents adopt a certain protocol; again, leading to a deadlock if all agents follow this strategy.

Transition versus steady state

Network externalities, timing of individual adoption decisions, and uncertainty can explain why the fact that a protocol is viable in steady state does not necessarily imply that there exists a transition path of adoption decisions leading to this state. In particular when regarding a succession of incremental innovations, there exists a minimum *innovation threshold* below which the *relative advantage* of the new protocol over the incumbent does not amortize the switching cost. This can cause protocols to fail even if they are technically superior to the incumbent. Examples of this phenomenon are plentiful.

Failure to reach the innovation threshold lets us rethink the very notion of failure. In fact, this kind of failure can be socially desirable if skipping an upgrade keeps a lower threshold for the adoption of the next incremental innovation that would not have met the higher bar. This observation dates back at least to work in the 1960s by Rogers [11], author of a sociology textbook on diffusion of innovations. (While diffusion is a broader concept spanning information flows, awareness and persuasion of people, it encompasses the stylized economic adoption decision in its core.)

The economists who studied network externalities in the 1980s added that the relative advantage depends on the number of adopters, leading to notions of *critical mass* (for the value of q at the tipping point) and game-theoretic models of possible interventions that solve the bootstrapping problem. Ozment and Schechter [9] have summarized these results for the case of security protocols including a case study on the adoption of DNSSEC.

Network topology and risk

Non-trivial graphical topologies underlying a network of externalities add substantial complexity and emergent features. The traditional "macro" view on network externalities implicitly assumes a fully connected graph between all participating agents: Metcalfe's law postulates that the value of a network of compatible communication devices is quadratic in the number of subscribers. This presumes that every *potential* communication relation adds the same unit of value. However, the objective functions of real agents are less symmetric, suggesting to model externalities on graphs. With real topologies hard to observe and even harder to tract, special classes of graphs (trees, bipartite, etc.) can be of interest to study typical phenomena. For example, the notion of *indirect* network externalities [5], where utility does not depend on the number of agents participating in the same but in a specific reference network, can be framed as feature of topology. (Think of complementary goods, such as payment system adoption depending on the number of merchants accepting it.)

Network topology is also important when externalities are negative, such as the propagation of risk between interdependent agents. Consideration of risk is a relatively young avenue in network economics [1], but the recent awareness of security in cyberspace may make it a more relevant success factor for the adoption of Internet protocols than ever.

2.2 Bitcoin

Bitcoin is a serious attempt to establish a global cryptographic currency in a fully decentralized manner. The concept was first described in a white paper published under pseudonym [8] and is now being developed by an open source community. The core of the Bitcoin system consists of a protocol, specified by reference implementation, and a global state stored in a distributed data structure called *block chain*.

Protocol

Technically, Bitcoin is a distributed peer-to-peer accounting system where account numbers are public keys. Account ownership is defined by knowledge of the private keys, which are used to sign transactions. Bitcoin uses majority consensus to enforce integrity such that the global state satisfies two constraints:

1. a non-negativity constraint for every account, thus preventing double-spending; and
2. an accounting identity for the sum of all account balances, thus ensuring conservation of value.

The latter distinguishes Bitcoin from earlier proposal of cryptographic cash. Replication ensures availability of the system state with high probability. A proof-of-work scheme discourages sybil attacks against the consensus mechanism. The protocol entangles the proof-of-work defense with an (approximately) incentive-compatible bootstrapping mechanism that regulates the initial distribution of wealth. This is known as *mining* process: solving cryptographic puzzles is rewarded with units of the virtual currency.

Ecosystem

Over the past two years, a vibrant ecosystem has developed around the Bitcoin core including merchants, exchanges, mining pools, remote wallets, and casinos. At the time of writing (late August 2013), blockchain.info, a statistical service, reports 11.5 million bitcoin in circulation, each trading for a market price of 125 US$, leading to an overall market capitalization of 1.5 billion US$. The Bitcoin network

records 50.000 transactions per day accumulating to a global state of 9 gigabyte. The computational effort spent on the proof-of-work puzzles is in the order of 500 terrahash per second. For comparison, a heavily cooled Intel Xeon CPU reaches up to 50 megahash per second and some GPUs are in the order of a few gigahash per second.

Bitcoin's success is somewhat surprising because it has established cryptographic payments against the backdrop of many failed attempts to launch forms of electronic cash in the 1990s and 2000s, awkward economics against dominant incumbents (chiefly PayPal and credit cards), speculative attacks and hitches at key players in the ecosystem (e.g., popular exchanges), adverse press and associations with crime (often justified), and therefore credible threats of government intervention with the potential of nullifying all deposits. In sum, it is probably fair to state that Bitcoin's starting position was much more difficult than the one of many Internet protocols. Yet it throve, and still survives.

3. BITCOIN AS A MODEL

What can we learn from Bitcoin? Clearly, the protocol design heeds all success factors recommended in RFC 5218. This is a starting point, but not always sufficient. In particular payment systems are subject to indirect network externalities where a critical mass of merchants has to accept a means of payment until buyers adopt; likewise merchants could wait until enough buyers are willing to pay with the new system. Solving this social coordination problem (or at least reaching critical mass) by means of a mandate was not an option for Bitcoin. The position of incumbents (essentially the whole established payments industry) is very strong and the sector as a whole is heavily regulated. As a result, mandates in favor of an unknown crypto currency were extremely unlikely. Quite the contrary: even progressive technology activists, such as the Electronic Frontier Foundation, explicitly dissociated themselves from Bitcoin.

In these circumstances, Bitcoin's strength arguably lasts on three factors. First, the built-in reward system for early adoptors. By contributing to the distributed transaction authentication, miners earn a predefined amount of bitcoins. This reward declines slowly and exponentially over time, which indicates its purposeful design as a bootstrapping vehicle. What is more, the difficulty of the mining puzzles is adjusted in a shorter control loop depending on the available network hash rate. Consequently, the probability of securing the fixed reward is indirectly proportional to the number of participants. This is a smart way to offset the barrier of entry imposed by network externalities. The author is not aware of any other protocol using this approach.[1]

The second success factor is not primarily in the protocol, but in the ecosystem. Observe that bootstrapping is inhibited by predominantly *indirect* network externalities whereas the reward mechanism is designed to offset *direct* network externalities. Bridges can be found in the ecosystem where third parties offer their services to exchange between bitcoin and conventional currencies. This allows early adoptors to interface with merchants accepting the incumbent payment networks. In other words, exchanges are institutions that convert indirect into direct network externalities. (Why direct? Because exchange fees and trading risk are indirect proportional to liquidity and thus decrease with a growing number of Bitcoin adopters.) In the language of network economics, exchanges serve as *adapters* which interface between two otherwise incompatible networks [4].

The third success factor lies in the interpretation of Bitcoin as money. Among other functions, money serves as store of value and thereby solves the inter-temporal matching problem in an exchange economy. Early adoptors are rewarded with the promise to exchange their bitcoins for something they desire at a future point in time. Trusting this promise, they may tolerate a limited number of merchants at the time of adoption in expectation of ongoing uptake by a wider set of market participants. Of course this is nothing else than speculating on the success of Bitcoin, which is a self-fulfilling prophecy; in particular if markets are incomplete so that building a short position on Bitcoin (i.e., speculating against it) involves high transaction costs.

So can we take Bitcoin's success as a model for other Internet protocols? Granted, few Internet protocols will have as difficult a starting position as Bitcoin. Those that struggle to get adopted nonetheless can take inspiration from Bitcoin's first two success factors, namely a built-in reward mechanism and an ecosystem providing adapters. The third success factor seems to be specific to protocols that create lasting virtual value (or comes with such a promise). This factor is not easily transferrable to arbitrary protocols.

4. BITCOIN AS ENABLER

If we cannot copy all of Bitcoin's success factors, can we at least use Bitcoin (or similar payment systems) to solve incentive issues in the design of other Internet protocols? Indeed, the canonical response to externalities in economics is to *internalize* them. This involves a transfer of value from agent j, who benefits from agent i's action, to agent i; or vice versa if agent i's action causes a loss for agent j. Typical network protocols lack the means for arranging such transfers. This constrains the design space of protocols to a mixture of self-enforcing, win-win, or dependence on altruistic action (i.e., some nodes absorb negative externalities).

In fact, many of the popular Internet protocols require that the systems implementing the protocols do not always operate in their owner's best interest. Of course, if large imbalances in the distribution of benefits and costs accumulate, the issue may be escalated to out-of-band mechanisms, such as the negotiation table. The debate on network neutrality is a vivid example, illustrating how hard it is to bargain a fair solution. With weak Internet governance and strong vested interests, a possible way forward is to provide means for internalizing the externalities right in the protocol design. We shall briefly explore three options.

Paying with money

Paying with money is typically *out-of-band*, informal, and in bulk volumes to keep the transaction costs low. As a consequence, this method requires a contractual relationship between identified legal entities, trusted metering and audits. And it locks out small players who cannot amortize the upfront cost to establish such relationships. Paying with money *in-band* is conceivable, but impractical. This is due to prohibitively high transaction costs of micro-payments. Piggybacking on a system like Bitcoin is a new avenue to explore. Possible caveats are computational overhead and latency of Bitcoin transactions as well as new security risks

[1]The author appreciates relevant pointers by the workshop participants.

as protocol stacks will need to know the private key, the digital counterpart of a blank check.

Paying with data

Some Internet services – think of them as informal protocols on the application layer – try to internalize externalities vis-a-vis end users by collecting and monetizing personal data. The typical exploitation strategy is to sell targeted ads. However, personal data is a very inconvenient unit of account. Personal data has all the disadvantages of information goods, it is subject to specific and complicated privacy laws, and it is very hard to measure the value of individual data points, which can be extremely volatile [2]. The law of large numbers tells us that only big organizations can balance the risk of receiving good and bad data points and extract benefits from a pool of personal data close to the (more predictable) average value. Moreover, paying with data is clearly not an option for protocols on lower layers.

Protocols with built-in payment systems

If our protocols could readily use direct value transfers between agents (such as money) instead of unwieldy substitutes (such as personal data or out-of-band settlements), it would be much easier to design incentive-compatible as well as privacy-friendly protocols and services. This suggests to design protocols with built-in payment systems, each lightweight and tailored to the purpose of the protocol. Researchers of peer-to-peer networks have made some steps in this direction to discourage free riders (a special form of externality) [6]. But is seems that more general classes of protocols can benefit from built-in payment systems. These systems can reward early adoptors and internalize the externalities that emerge during operation alike. As more protocols with payment systems emerge, third parties can offer exchange services between different protocol currencies. An advantage of clearing in a specific currency rather than in bitcoin or dollar is that security and privacy measures can be tailored to the application, making it harder to loot nodes or launder money, but admittedly not impossible. Recall that many protocols do not have effective defenses against malicious resource consumption (denial of service) and rely on monitoring and detection to prevent abuse.

Note that it will not always be easy to exactly quantify an externality. Instead, we have to devise approximations, possibly refined with feedback loops or simple market mechanisms for adjustment at runtime. Research on approximate mechanism design can help to gauge how close we can get.

5. OUTLOOK

We have explained critical success factors for protocol adoption in terms of network economics. The most salient feature of the Bitcoin protocol is its reward mechanism for early adoptors. Although not all ingredients of Bitcoin's success generalize to arbitrary protocols, having built-in payment systems could substantially widen the design space for easily adoptable protocols. Payments enable mechanism to internalize the externalities that emerge during the transition phase as well as in steady state.

Protocols with built-in payment systems call for a larger research agenda. Technically, how can payment functionality be built in a reusable manner so that protocol designers, who are not necessarily specialists in payment systems security and privacy, can easily embed them in their protocols? How to standardize mechanism designs? How to secure funds in the protocol stack through system events such as virtualization, cloning, or reboot after a crash? What about regulatory implications such as taxation, financial regulation, and law enforcement?

Besides the idea of built-in payment systems, there are open research questions on protocol adoption. For example, can we empirically validate and rank the importance of the success factors in RFC 5218? If altruism was an important driver of early Internet development, whether and why did it disappear? Is this primarily related to scale (leaving more money on the table is less tolerable), intensified competition (every cent counts), or better information (knowing how much my partner takes away makes me eager to get my share)? Technical records of protocol adoption, if available, can be valuable resources for interdisciplinary research.

A final remark concerns the hypothesis that Bitcoin's subversive nature is a success factor in itself, an argument that can also be made for systems like Tor or BitTorrent. We have no evidence to tell whether a fraction of Bitcoin adopters with illicit goals – chiefly money laundering [7] – had just no alternative and thus brought the critical mass to unleash sustained adoption. If there is some truth in it, this is certainly a factor that should not make it into an RFC.

6. REFERENCES

[1] D. Acemoglu, A. Malekian, and A. Ozdaglar (2013): *Network Security and Contagion.* Working Paper. MIT Department of Economics.

[2] S. Berthold and R. Böhme (2009): Valuating Privacy with Option Pricing Theory. Eighth Workshop on the Economics of Information Security (WEIS), University College London, UK.

[3] B. Briscoe, A. Odlyzko, and B. Tilly (2006): Metcalfe's Law is Wrong. *IEEE Spectrum* **43** (7) 26–31.

[4] N. Economides (1996): The Economics of Networks. *International Journal of Industrial Organizations* **14** (2) 673–699.

[5] M. Katz and C. Shapiro (1985): Network Externalities, Competition and Compatibility. *American Economic Review* **75** (3) 424–440.

[6] D. Levin, A. Schulman, K. LaCurts, N. Spring, and B. Bhattacharjee (2011): Making Currency Inexpensive with iOwe. Sixth ACM NetEcon Workshop, San Jose.

[7] M. Möser, R. Böhme, and D. Breuker (2013): An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem. APWG E-Crime Researchers Summit.

[8] S. Nakamoto (2008): *Bitcoin: A Peer-to-Peer Electronic Cash System.* `http://bitcoin.org/bitcoin.pdf`.

[9] A. Ozment and S. Schechter (2006): Bootstrapping the Adoption of Internet Security Protocols. Fifth Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK.

[10] D. Thaler and B. Aboba (2008): What Makes for a Successful Protocol? RFC 5218.

[11] E. M. Rogers (1962): *Diffusion of Innovations.* Free Press.

[12] C. Shapiro and H. Varian (1998): *Information Rules.* Harvard Business School Press.