

Autonomics

By Joannes Vermorel, CEO of Lokad, May 29th, 2018

Abstract: We propose the definition of a new field of study named *autonomics*. This field is characterized by the study of complex machine-driven systems purposefully built around well-chosen economic incentives which ensure the survival of those systems. We argue that *Bitcoin* (Satoshi Nakamoto, 2008) was a foundational study that triggered the emergence of autonomics. The introduction of a distinct field of study is motivated by the fundamental misunderstandings that emerge when systems - like *Bitcoin* - are approached from more traditional perspectives. Beyond this, autonomics may provide a superior perspective to reason properly about Bitcoin, its peers and its dependent sub-systems.

Overview

Within the 10 years since its publication, *Bitcoin* (Satoshi Nakamoto, 2008) has led to the emergence of a new class of economic systems which appear to have surpassed 500 billion USD in market capitalization at some point in December 2017¹. From an epistemological perspective, back in 2008, Bitcoin was a remarkably alien intellectual construction which largely baffled many prominent figures of both the academia and the business world alike. Even more puzzling, even among *believer's circles*², is that Bitcoin has been routinely so thoroughly misunderstood that this very system has been routinely damaged to the point it was not even working anymore as *digital cash* which was unambiguously the whole point of the original publication by Satoshi Nakamoto.

We propose *autonomics* as a new field of study for Bitcoin and its cousins. The term is coined from the contraction of *automation* and *economics*. We argue that analyzing those systems through the lenses of autonomics is a requirement if one wants to ensure that those systems remain both functional from a computing perspective, and sustainable from an economic perspective.

Furthermore, we argue that, lacking for a perspective superior to *autonomics*, attempts at analyzing those systems from the perspective of alternative - more established - fields of study

¹ As those markets are still nascent, assessing reliable valuations remains difficult. This complexity is compounded by the level of volatility associated with those markets as well. Nevertheless, causal observations of the state of affairs of the world in 2018 indicate that *Bitcoin* and its competitors have become part of the global economic landscape.

² A loose characterization of the groups of people that have a highly positive opinion about Bitcoin or similar systems, irrespective of any particular understanding of those systems. Depending on the degree of their belief, *believers* may also invest a portion of their personal assets in those systems, primarily expecting their wealth to grow through the growth of those systems.

are misguided. Worse, it may give an aura of authority to authors who should not be considered as competent peers in autonomics *only* due to their contributions in alternative fields³.

A first definition of autonomics

Autonomics is the study of a specific class of complex systems. The complexity refers to the size of the population - expected to be large - of people who socially interact with the system through machines and their user interfaces. Moreover the system has the following properties:

- **unified:** the boundary of the system is defined by its purpose. Any part that is required to keep the system working properly should be treated as being part of the system.
- **sustainable:** the boundary of the system has to be defined in such a way that it includes all parts that make the system economically viable.
- **automated:** the system operates without any arbitrary caps on correctness or diligence. No system that relies on human interventions to operate can meet this criterion.
- **permissionless:** the only entry barrier to the system is competency at conducting operations within the system. The system internally behaves like a market, funding its own sustainability.

The **unified** property refers to the purpose (i.e. the *function*) of the system which is properly identified and understood as such by the participants. This property characterizes the *boundaries* of the system. Indeed, *complex systems* are somewhat elusive, and as everything tends to be interconnected in modern economies, it's usually difficult to establish the boundaries of the system under study. Failing at properly defining the boundaries of the system leads to incorrect reasoning about the system itself. We propose that the system is best defined as the sum of all the elements contributing to the resolution of the same problem. In Bitcoin, the system is unified by its intent to be *_electronic cash_* as stated in the title of the original Bitcoin paper.

The **sustainable** property indicates that the system as a whole benefits from the economic equivalent of an autocatalytic chemical reaction - self maintained until its reagents are exhausted: The system keeps operating until the willingness to pay of the participants is exhausted. The system needs maintenance, as many machines are involved, but the maintenance is paid for by the participants who extract value from the ongoing operating state of the system. This property further refines the boundary: All the parts that are required to achieve economic viability have to be considered as part of the system. For example in Bitcoin, sustainability is ensured by two complementary mechanisms: first the block emission rewards and second the transaction fees.

³ As autonomics lies at the intersection of several fields of study, common sense dictates that being a recognized authority in any of those alternative fields is a positive element when it comes to make statements about autonomics. We are merely arguing that this is not a *sufficient* condition.

The **automated** property indicates that the system fulfils its purpose with a degree of correctness and diligence that is incompatible with what humans can deliver without computers. This property narrows the scope of *autonomics* to systems characterized by their use of computing resources at a foundational level, to the point that humans are essentially superfluous for all routine operations. Humans are only involved when it comes to the *upgrade* or the *maintenance* of the system. However, by virtue of the sustainability property, the work performed by those humans is also paid for by the wealth generated by the system itself. At the present time, due to technological limitations, the only complex tasks where machines can operate with a degree of automation compatible with the scope of autonomics are informational tasks⁴.

The **permissionless** property indicates that any participant can join the system with no *superfluous* entry barrier except her own competency at operating within the system. This property narrows the scope to systems that internally behave with market dynamics. This property confers antifragility to the system, as markets operate as autonomous filters on the competency of participants. Yet, in any system - more generally any market - has entry barriers as there is always an element of capital needed to get started⁵. We are referring here to artificial elements that would arbitrarily raise the cost of entering the system without efficiently generating a corresponding gain⁶ for the system.

In the following, the term *system* always refers to the autonomics perspective. Hence, a system is expected to exhibit the four properties listed above. Those systems can be further characterized by a cursory analysis of the consequences of those properties.

A few corollaries

The characterization of the systems that are of interest from the perspective of *autonomics* offers a series of immediate corollaries. Those corollaries are of interest to gain a better grasp of the nature of autonomics and they offer some insight on future evolutions of those systems.

⁴ While this proposition is highly speculative, technologies like additive manufacturing, robotics and machine learning might progress to the point where systems oriented toward *material purposes* (eg. keeping streets clean through robots printed and maintained on demand) rather than *informational purposes* (eg. money) might become both technically feasible and economically viable. If or when this becomes possible, the autonomics field of study will have to be revisited with two branches: first *soft autonomics* dedicated to informational purposes, second *hard autonomics* dedicated to material purposes.

⁵ The acquisition of competency requires some upfront capital, and, in practice, a bit of luck as well. Thus, no system is ever absolutely permissionless. Nevertheless, it would be very incorrect to use this argument to conclude that it's a continuum from permissionless to centralized, as there are strong economic nonlinearities when going from one end of the spectrum to the other.

⁶ Mining in Bitcoin, i.e. securing enough hashrate, to have a chance to successfully win one block per year, comes to represent now a rather substantial capital investment. However, this does not contradict the permissionless property because winning the hashrate game is foremost a matter of superior execution, which can be profitably executed at almost any scale; i.e. a participant can start outcompeting its peers at a small scale, and gradually grow into a system leader.

The market validates the systems: The combination of the sustainable and permissionless properties implies that the correctness of any proposition made in the field of autonomics can and should be assessed from an economic perspective, measuring how good the proposition fares in the world at large⁷. Unlike general economics which typically cannot be experimented upon⁸, autonomics are a highly empirical field of study. Anyone that claims having a superior proposition can start building an allegedly superior system and see whether it outcompetes others. Markets do exhibit short-term volatility, but any proposition that is contradicted by market behaviors for a long duration should be rewarded as incorrect.

Bots cannot be denied access to the system: The automation property enforces that all the routine operations of the system should be accessible to a pure mechanical entity, commonly referred as a *bot* in software jargon. Indeed, if the parts of the system cannot be operated by a bot but require a human intervention instead, then the system cannot exhibit the correctness and diligence that are required to be considered a valid candidate for autonomics. In particular, this implies that any system that requires any kind of "stamp" from a bureaucratic entity cannot be considered part of autonomics unless the bureaucratic entity itself is entirely automated.

The system can only be distributed: The permissionless property is not compatible with any system except the distributed ones. In computer science, *distributed computing* refers to calculations that are operated on many machines, which is also referred to as horizontal scalability. Traditionally, distributed computing involves fleets of machines that collaborate to fulfill a specific function by pooling their respective computing resources for the task. In the case of autonomics, the purpose of the distribution is not necessarily to pool computing resources, it might just be there to achieve a trust-and-verify scheme where every participant monitors a portion of the state of the system - and possibly the entire state. Then, depending on the design of the system, each participant may require some horizontal scalability for her own setup as well, although this latter aspect is not a requirement for autonomics.

The system can only be scalable: The permissionless property requires systems to scale up to mankind-scale. Indeed, any system that comes with built-in quotas on essential properties⁹,

⁷ From an autonomics perspective, the strongest element in favor of the selfish miner hypothesis proposed by *Majority is not Enough: Bitcoin Mining is Vulnerable*, November 2013 (Ittay Eyal, Emin Gun Sirer), <https://arxiv.org/abs/1311.0243>, is not the internal correctness of the analysis, but the fact that the selfish miner hypothesis has now been proven true through the Monacoin attack in May 2018 (<https://www.ccn.com/japanese-cryptocurrency-monacoin-hit-by-selfish-mining-attack/>). Nevertheless, the success of the attack does not imply that the Nakamoto consensus is terminally flawed, merely that a specific attack vector - selfish mining - needs to be taken into account.

⁸ States have distinct economic policies which can be benchmarked. However, the process is extremely slow and the number of data points that result from the economic "experiments" is too low in practice to be able to draw any conclusion beyond some general principles (e.g. quotas make markets dysfunctional) - most of them having been known for more than a century.

⁹ In Bitcoin, it's not possible to have a single transaction that moves money from or to 1,000,000 endpoints at once: There is a limit on the maximal size of a given transaction. However, this limit should

scalability being one of them, is denying the permissionless property to the system, which, in turns, breaks all the internal market dynamics of the system. If a system can only operate at mankind-scale but only through parts that are deemed external to the system, then the boundaries of the system are improperly defined: Those parts should be considered as being an integral part of the system¹⁰.

Systems cannot and should not be made technology-resistant: As a system is expected to be both automated and sustainable, it fundamentally seeks to outcompete other systems through better technology and better incentives. Thus, any attempt at making a given system resistant to a particular technology, such as ASIC, is economical nonsense. In autonomics, systems do not try to resist newer technologies, they embrace them fully to outcompete alternatives.

Systems are neutral with respect to regulations: If joining the system requires a governmental approval then the system is not part of autonomics as it violates the permissionless property, and probably the automation property as well - unless the government provides fully automated communication endpoints to support the process. Yet, while in autonomics, systems are neutral - by essence - when it comes to regulations, participants are not, and remain bound by their respective applicable regulations. Conversely, no government can be denied access to a system by virtue of the permissionless property.

Systems also compete on education: The prime entry barrier to the system is the competency for a participant to join the system. Depending on the intent of the participant, the required amount of competency ranges from almost trivial (e.g. using a simple app on a mobile) to extremely challenging (e.g. designing a superior class of ASICs that outcompetes, power-wise, existing designs). As systems grow stronger - economically - through their own momentum, they also compete on educating the world at large to spread the competency at being operated. Thus, it can be expected that highly competitive systems will invest in out-educating their competitors.

While those corollaries are fairly immediate from the perspective of autonomics, it would have been difficult to reach the same conclusions from the perspective of more traditional fields of research.

not be interpreted as a quota because it does not prevent any relevant use case from happening as far as *electronic cash* is concerned.

¹⁰ Bitcoin Core cannot be considered as a relevant system from an autonomics perspective as the quota on the block size prevents new participants from joining the system, which denies the permissionless property. The jury is still out to figure out whether Lightning Networks can be considered as proper autonomics. However, the inner economical dynamics within the Lightning Networks strongly hint that the first-mover advantage conferred to payment hubs will most likely result in a stable oligopoly denying its permissionless property in the future.

A novel field of study

We propose that systems like Bitcoin cannot be properly understood and analyzed from the perspective offered by the more traditional fields of study. More specifically, we claim that framing autonomics through the lenses of those alternative fields of research leads to conclusions that are factually incorrect, and typically proven wrong by the market itself. In the following, we review a list of perspectives that we believe to be insufficient to characterize such systems.

Cryptography: While asymmetric cryptography appears to be a practical requirement to build systems, it predates Bitcoin by more than 3 decades as the RSA was originally patented in 1977. Similarly, one-way hashing functions - which are used to perform the proof of work - have been around for even longer, as their history can be traced back to the 1950s at IBM. Systems like Bitcoin make clever use of cryptography just like they also make clever use of a P2P network. Yet, cryptography is not a defining trait for those systems, merely a technicality. Naturally, superior cryptographic primitives are desirable. However, as the cryptographic primitives presently in use¹¹ are close to theoretical limits, it is unlikely that cryptography will be a fundamental differentiator¹² between systems. Trying to analyze and improve a system only through the lenses of cryptography leads to nonsensical decisions that endanger the sustainability of the system, for example by failing at taking the appropriate course of action to keep the system properly operating so that its participants can extract the desired value from the system¹³.

Economics: Claims in economics can and should be validated by the corresponding observations in the markets themselves. However, the general market dynamics do not lend themselves easily to empirical validations. Economics are fundamentally a matter of observation, and of establishing principles that seem to most generally guide markets and their participants. This perspective is quite different from the one of autonomics which aims at providing the necessary insights to build and maintain systems. Autonomics comes with a dose of engineering that is profoundly absent from traditional economics. Even sub-fields such as *computational economics* are merely incremental variations on traditional economics rather than radical re-foundations. As a matter of anecdotal evidence, during the first decade of existence of autonomics, mainstream economists have been largely absent from the scene, and none of them seem to have made any notable contribution to any system so far.

¹¹ Bitcoin is already using efficient cryptographic primitives. It can be expected that systems will compete on reliability, usability and overall mindshare beyond just the security angle.

¹² If a cryptographic primitive is discovered to be broken, it will be swiftly replaced by another one that is not. Such an event will disrupt the system as it would immediately trigger the need to perform maintenance operations at many levels. However, if the maintenance is properly executed, long term effects on the system should remain minimal.

¹³ Bitcoin Core, one of the notable forks of Bitcoin, failed at keeping transaction fees under control, which, in turn, denied the possibility to the participants to extract value from the Bitcoin Core network.

Distributed computing: This branch of computer science is relevant to the design of autonomic systems as the permissionless property imposes a distributed design to make it possible for further participants to join. For example, Bitcoin adopts a small world peer-to-peer distribution design, but other approaches are conceivable. Then, as a system needs to scale, some of its parts may require some horizontal scalability, which is also a matter of distributed computing. For example, the UTXO dataset needs to be maintained at scale, which is a small variation of the well-known *key-value store* problem. Yet, while distributed computing is quite similar to autonomics due to the amount of engineering involved, the sustainability perspective is largely absent: In distributed computing, solutions are expected to be maximally efficient, in terms of computing resources, to address a given problem. However, solutions are not expected to fund their own market survival.

The discussion presented in this section sheds light on why traditional fields of study do not offer satisfying perspectives when it comes to the study of the systems of interest for autonomics. For the sake of brevity, we are only reviewing the short list of relevant fields although *software engineering, politics, hardware design, education, psychology* or *UX design* would be equally relevant fields of study which provide invaluable insights for autonomics.

Conclusion

The brief history of Bitcoin indicates that the sustainability of a given system should not be taken for granted. Security flaws in the original design of the system can be exploited by adversaries, and do require prompt corrective actions. Minor imperfections can gradually turn into massive problems as the system grows, and the fix may require substantial re-engineering of the system itself. Also, as technology progresses, a system can find itself threatened by competing systems which are leveraging better designs. Remaining competitive may involve even greater re-engineering for a given system. Thus, all systems require maintenance and, over time, some degree of inevitable change as well.

Yet, the improper maintenance of a system can lead to the undesirable outcome of making the system a less valuable option for the very problem it tries to address. For example, the misguided maintenance of a notable fork of Bitcoin (Bitcoin Core) led to an explosion of transaction fees in 2017 which was preventable as notable voices had anticipated this very problem years in advance. While the root causes of this improper maintenance are complex, the lack of an adequate perspective - as proposed by autonomics - to even frame the problem was a clear handicap to form a community consensus for the resolution of the problem.

The ambition of autonomics is to foster a field of research that will generate insights that are most relevant to the maintenance and to the ongoing improvement of a system. While politics are unavoidable whenever a large number of humans is involved, systems that ground their own maintenance in rational findings will outperform those that do not. Autonomics introduces entirely new classes of failure modes. With the explosion of digital currencies in the last few

years, many newer failure modes have already been observed in the market. Some other failure modes remain theoretical, possibly because of proactive actions from most system maintainers.

The long term sustainability of a system depends on the ability to pass the knowledge required to perform its maintenance from a generation of maintainers to the next. As a body of knowledge, autonomics intends to make this possible.