

# There is plenty of time at the bottom

## The economics, risk and ethics of time compression

Anders Sandberg, Future of Humanity Institute, Oxford Martin School, University of Oxford

### Abstract

**Purpose** – The speed of computing and other automated processes play an important role in how the world functions by causing “time compression”. This paper reviews reasons to believe computation will continue to become faster in the future, the economic consequences of speedups, and how these affect risk, ethics and governance.

**Design/methodology/approach** - Brief review of science and trends followed by an analysis of consequences.

**Findings** – Current computation is far from the physical limits in terms of processing speed. Algorithmic improvements may be equally powerful but cannot easily be predicted or bounded. Communication and sensing is already at the physical speed limits, although improvements in bandwidth will likely be significant. The value in these speedups lie in productivity gains, timeliness, early arrival of results, and cybernetic feedback shifts. However, time compression can lead to loss of control due to inability to track fast change, emergent or systemic risk, and asynchrony. Speedups can also exacerbate inequalities between different agents and reduce safety if there are competitive pressures. Fast decisions are potentially not better decisions since they may be made on little data.

**Social implications** - The impact on society and the challenge to governance is likely to be profound, requiring adapting new methods for managing fast-moving and technological risks.

**Originality/value** – The speed with which events happen is an important aspect of foresight, not just as a subject of prediction or analysis, but also as a driver of the kinds of dynamics that are possible.

**Keywords:** time, computation speed, limits of computing, automation, productivity, timeliness, competition, algocracy, decision speed, systemic risk

**Article classification:** Conceptual paper

### Introduction

Whether there is a long-term trend towards accelerated change is controversial (Kurzweil 2010; Cowen 2011; Eden et al. 2012; Sandberg 2013), but clearly the present era is experiencing a remarkable increase in the speed of computation, a compression of the time required to perform a computational task. This paper is an exploration of the limits and implications of this compression on human values and society.

If everything speeded up equally, there would be no change: one of the peculiarities of time (and a big source of arguments between Platonists and Aristotelians in the philosophy of time) is that we mostly notice it through the change of events and things relative to each other (Markosian 2016). “Speeding up” hence means that more things of a certain kind occur relative to other things, or that a kind of event occurs before another event it would previously have occurred after.

Faster computation means that computational goods can be produced faster and earlier. This paper will explore some of the consequences and limits of this phenomenon.

What kinds of value would accrue from something occurring faster?

- More instances of the event in a given interval: if these hold value, then there is more value produced. A faster manufacturing process will produce more goods per unit of time.
- Time is freed up by the faster rate of work, and this is valuable. For example, a labour saving device frees up time that could be used for leisure or more productive work.
- Having an event occur earlier than another event, becoming valuable because of the ordering. For example, diagnosing and intervening rapidly against a medical condition before it turns serious.
- The value of a remote event is increased (possibly from zero) by it becoming closer to the present. This includes becoming able to achieve something that previously was impossible to achieve within the given timeframe. For example, calculating weather forecasts or solutions to mathematical problems that previously would have taken years in hours.

Limits to speeding up computation include physical limits, but also limits due to the difficulty of tasks (or the algorithm used to solve them). At present we appear to be far from any fundamental technical limit on computing power, but we have already touched the fundamental limit on communication speed. Speeding up the interaction with the physical world may prove challenging because of the discrete nature of signals and sluggish responses of macroscale actuators.

Faster computation does raise risks and ethical challenges: various forms of loss of control, inequalities of speed, gaps between oversight and system speed, loss of opportunity due too early decisions, and possibly so much change that the “change budget” becomes depleted. In particular, speedups appear to pose a serious challenge to human ability to control technological processes due to growing gaps of speed between computation and control (“cybernetic gaps”) and challenges to setting the goals they are optimizing for due to gaps of speed between computation and the human world (“ethical gaps”), in turn posing a profound challenge to governance systems that are themselves to some extent hybrid human-computational systems suffering internal speed gaps.

## Limits to computation

Computation is subject to both physical, logical and technical limits. This section will outline some of the main limits and how they limit achievable computational speedups.

### *Physical limits to computation*

Computation requires structured change in information-bearing material systems, making the speed a particular computation can be performed at limited by the speed in which a physical system can change in a corresponding way.

The most fundamental limits are due to quantum speed limits, stating how fast a quantum system can move between two distinguishable states. The Margolus-Levitin limit states that a system with mean

energy  $E$  cannot move to another orthogonal state in less than  $\pi\hbar/2E$  time, and the Mandelstam-Tamm limit  $\pi\hbar/2\Delta E$  where  $\Delta E$  is the standard deviation of the energy of the system (wrt to initial state). Later work has found entire families of quantum speed limits, where the bound scales as  $1/\Delta E$  for unitary dynamics (Margolus & Levitin 1998; Pires et al. 2016).

These limits, together with the Bekenstein (1981) bound on the information that can be contained within a region with given radius and mass-energy, can be used to describe the most extreme computer systems that could be built even in principle: for a one kilogram, one litre computer the limit is  $5.4258 \cdot 10^{50}$  logical operations per second on  $\approx 10^{31}$  bits (Lloyd 2000). In this case the state of mass-energy is mostly akin to a black hole or a small piece of the Big Bang.

The speed of computation in normal matter is limited by the number of transitions the system can perform per unit time without breaking down. This in general depends on the available energy to perform the transitions and how strong the energy barriers of the system are. Semiconductor and optical devices can switch on the picosecond scale (Mii et al. 1994; Ctistis et al. 2011). Molecular computation is limited by bond energies: above  $10^{15}$  transitions per second the energy involved becomes larger than the bond energies and the system starts to break up (Drexler 1992). In terms of switching current computers are hence about five orders of magnitude slower than the hard limits unless nuclear matter computation on the yoctosecond scale eventually becomes possible (Sandberg 1999).

Very fast computation is also very small and localized, since the light speed limit forces the parts involved to be closer than  $t/c$ , where  $t$  is the cycle time. A nanosecond is about a foot long, a femtosecond 300 nanometers. Since there is also a limit to how densely information can be packed (likely on the order of a bit per atom for molecular systems) this means that a computation taking time  $t$  cannot process more than  $4\pi\rho(t/c)^{1/3}/3$  bits. If individual atoms perform computations, the cycle time they would require to exchange state information at lightspeed in diamond would be  $1.2 \cdot 10^{-18}$  seconds.

Quantum computation does not change this fundamental issue. The number of steps an algorithm needs to perform in order to arrive at a solution of a problem defines the complexity class of the problem. A practically useful algorithm scales polynomially or less with problem size. Quantum computing merely leads to an exponential speedup of certain problems, which makes *some* computations that would otherwise be infeasible potentially doable (given quantum computers) (Moore & Mertens 2011). It should be recognized that innovations in algorithms can also make classical computations significantly faster: the FFT algorithm changed the complexity of discrete Fourier transforms from  $O(N^2)$  to  $O(N \log N)$ , making e.g. online multimedia feasible.

### *Algorithmic limits*

The physical limits discussed above represent limits on the speed of an algorithmic step: the complexity classes represent limits of how fast problems can be solved. For example, it is known that any sorting algorithm that must compare elements with each other has to run in  $O(N \log N)$  time on one processor, and if it is parallelized  $O(\log N)$  (Moore & Mertens 2011). Any improvement beyond this will have to come from refactoring the problem (e.g. making use of known constraints on the data, such as integer sorting).

There is a key difference between parallelizable tasks and serial tasks. By splitting a problem into suitable parts a solution can often be generated far faster on more processors. For example, summing  $N$  numbers requires time proportional to  $N$  on one processor, but can be done in  $\log(N)$  time on  $N$  processors (first every pair of processors sum their numbers, then the  $N/2$  partial answers are summed, and so on until only one answer remains). Other problems have data dependencies that make this impossible, for example calculating trajectories in the 3-body problem. The key difference is whether the critical path length  $C$  of the computational dependency graph is close to the overall amount of computation  $T$ . If  $C$  grows more slowly than  $T$ , then parallelism gives a time gain. However,  $C$  represents an insurmountable problem-dependent barrier for how fast the problem can be solved (this is the cause of “Amdahl’s law” (Amdahl 1967) limiting the speedup of a task as more processors are added).

While the complexity classes represents the logical bedrock of how fast tasks can be done, one should not underestimate the potential speedups due to refactoring, approximating, or “cheating” on problem instances that matter. For example, the 3-SAT problem is known to be NP-complete and should hence not be expected to have efficient solutions, yet heuristic 3-SAT solvers are successfully used in circuit design and automatic theorem proving. They are not efficient on all possible instances but for practical purposes they work (Moore & Mertens 2011).

It is hence harder in general to estimate the distance to algorithmic limits to computation than physical limits because real-world problems are rarely well specified enough. Improvement in algorithms are often 50-100% of the gains from hardware progress (Grace 2013). In addition, predicting future theoretical insights is often as hard as gaining the insight itself: strongly idea-driven technological change is by its nature less predictable than incremental change.

## Technical limits




Figure 1: Distribution of sigmoid curve scenarios of computer power generated from data in (Koh & Magee 2006). The grey shade denotes density of scenarios. The red curve is the median scenario, green lines the 5% and 95% percentile scenarios.

Moore's law has been formulated in numerous incompatible ways (Mack 2015; Waldrop 2016), but perhaps the most relevant measure of progress is processing operations per second per dollar. Merely measuring speed will not capture the actual practical impact. This measure has been growing exponentially over several decades and even if one fits a pessimistic sigmoid curve (implying that growth must eventually come to an end) the median estimate implies about 20 orders of magnitude improvement (!), and with 95% probability at least a factor of 100,000 improvement.

Moore's law is a self-fulfilling industry prophecy, partially driven by Wrightian learning and increasing production (Nagy et al 2013), partially by expectation (Schaller 1997). Given the value of faster computation there is a demand for it, and production makes new computational tasks feasible and affordable ("Bell's law") (Denning & Lewis 2016). Eventually it will run into physical limits (Mack 2015; Waldrop 2016), but compared to the previous section it is clear that there is a fair distance to run. Even if the technology itself stops the economies of scale may keep on increasing the performance per dollar.

## *Communications*

Communication has essentially reached the ultimate limit, lightspeed transmission. In electric cables there is some slowing (50%-99%*c*) due to electrical inductance and capacitance. Optical fibres typically transmit signals at 70% of lightspeed due to the refractive index. In contrast radio waves move nearly at the speed of light, but require line-of-sight.

Bandwidth is still increasing: Nielsen's law of Internet bandwidth suggests a 50% increase per year for users (Nielsen 1998). This is slightly slower than Moore's law, making accumulating data rational since it can be generated faster than it can be transmitted. The upper limits of bandwidth are set by the entropy of electromagnetic radiation, scaling as  $1.1194 \times 10^{21} \sqrt{A/d} P^{3/4}$  bits per second for an area  $A$  transmitter and receiver at distance  $d$  and using  $P$  Watt power (Lachmann et al 2004). We are clearly far from the physical limits yet, but were Nielsen's law to continue we will reach them before the end of the 21<sup>st</sup> century.

As mentioned above, lightspeed limits imply a space-time trade-off. A return time of  $t$  from sending a signal and getting an answer implies a distance  $D < t/c$ . Faster processing require smaller spatial systems.

Large systems will have parts that are causally disconnected: they cannot interact over processing cycles. Krugman's "First fundamental theorem of interstellar trade" applies here: interest rates for local information and for information in transit are the same (and by the second theorem, arbitrage will equalize them in different parts of the system). However, prices in different parts of the system may not be equal (Krugman 2010). While intended for space this also applies to a fast desynchronized computing world.

## *Sensing and acting*

To fast sensors the world is dim and noisy: the rate with which photons, sound or other measured entities arrive is slow and natural irregularities will dominate. If the intensity is  $I$  and sampled at frequency  $f$ , during each sample interval only  $I/f$  units of energy will arrive, an amount decreasing with  $f$ . If this interval becomes comparable to the average arrival time of measurable entities typically aliasing effects (if the arrival times are regular) or Poisson noise (if the arrival times are random) show up, drowning the signal in noise. To function well fast sensory systems need intense, high-frequency signals or sensors with a broad sensitivity.

The world is sluggish and hard to coordinate for fast controllers since the response from actuators will be slow and delayed. The time from sending a command until a response is received measured in processor cycles is  $2L/vf$ , where  $L$  is the size of the actuator,  $v$  the signal speed and  $f$  the processor frequency. This estimate optimistically assumes an instant response from the actuator, but for many physical systems response times are proportional to  $L$  (Drexler 1992), increasing the time to  $(2/v + K)L/f$  where  $K$  is the actuator sluggishness. The smaller and faster the actuators are the better the system can work, but this also requires closeness in space.

Fast computation hence benefits small systems acting in intense environments more than large systems dealing with uncertain, weak signals.

## *Summary*

In summary, we have good reasons to expect computing to become many orders of magnitude faster in the future: there is still plenty of distance to physical limits, and algorithmic improvement, innovation and (quantum) parallelization is possible in many domains. Indeed, “there is plenty of time at the bottom”.

At the same time we are close to communications and sensing limits, the improvement speed may be unpredictable, and it is hard to synchronize fast distributed systems.

## **The value of fast computation**

In the previous section, we discussed the physical limits to computation - how fast computation might be made in principle, if sufficient effort was dedicated to making it faster. For such effort to be expended, fast computation needs to have enough value to justify the investments in it. This section will review the reasons for why “faster is better” in computation, and how these reasons act as drivers for computational speedups and hence incentivize (due to economic profitability) further speedups.

## *Productivity*

If operations can be performed faster, then more instances can occur in a given interval. If these hold value, then there is more value produced. A faster manufacturing process will produce more goods per unit of time, a faster surgeon can operate on more patients. This is the normal world of economic productivity, endlessly studied by economists. Whether the value increases proportional to the speedup or not depends on whether the increased supply is enough to decrease price noticeably.

A closely related improvement is that time is freed up by the faster rate of work, and this is valuable. For example, a labour saving device frees up time that could be used for leisure or (more commonly) further productive work.

An effect of this is that alternative cost of time increase. A given time interval can be used for many more things, some of which are valuable. Hence wasted time may paradoxically become more of a serious problem in some domains. This might be a current contributor to information stress – there is always something of value going on, and it might be rational to switch tasks often in order to find high-value tasks. As described by Schwartz (2004) “missed opportunities” are often experienced as stressful since we notice lost utility, especially if the aggregated alternatives not taken sum to more utility than our choice. The actual rationality depends on how the overhead of searching is compared to the expected gain<sup>1</sup>;

---

<sup>1</sup> If different tasks have utilities per unit of time distributed as  $f(u)$  for some probability distribution  $f$ , then the best utility after sampling  $N$  will have distribution  $f^*(u) = Nf(u)F^{N-1}(u)$ , where  $F(u)$  is the

information foraging theory predicts that as the cost of switching between information sources decreases the time spent on each will decrease (Pirolli & Card 1999).

Changes in productivity will also lead to changes in allocation of resources (labour, computing). A sudden change in speed of one area will produce transient response in the labour/computing market before a new equilibrium establishes itself.

## *Timeliness*

Time is a non-renewable resource. Or rather, temporal location is non-renewable: what matters is often that X occurs before Y occurs. Since time is irreversible ordering effects can matter significantly, as in diagnosing a disease before it becomes life-threatening or inventing an offensive technology before any defence to it exists.

The Japanese earthquake early warning system uses seismological detectors to shut down trains before the earthquake wave hits (Kamigaichi et al. 2009) – once detection, transmission and reaction is fast enough this becomes possible.

The timescales in cars form another useful set of examples of how timeliness enables sharp shifts in performance. The trip planning timescale for humans is on the order of minutes, allowing it to be done automatically with 1990s technology; faster is more convenient, but once the navigation system passed the minute threshold it was good enough. The driving timescale involves decision-making on the order of tens of seconds to seconds. There was an important shift in 2006 when computing got fast enough, suddenly enabling autonomous cars in natural environments. Again there is a threshold effect: speedups of computing may enable better driving quality, but the fundamental ability to drive showed a fairly discrete transition. Similarly airbag systems gather sensor information and the control unit decides whether to trigger the airbag within 15 to 30 milliseconds after a crash has begun. This could be done already in the early of 1970s by dedicated electronics, and since then speedups have mostly served to make the decision-making more sophisticated (although improvements in sensors may have played a larger role). In all these examples computational speedups enable a new level of automation when they pass a human or mechanical speed threshold.

---

CDF. The median utility for the best is  $F^{-1}(2^{-\frac{1}{N}})$  and the expected utility is  $E[u_*] = N \int u f(u) F^{N-1}(u) du$ . In the case of Gauss-distributed utilities  $E[u_*] \approx \sigma \sqrt{2 \ln(N)}$ . The expected total utility of trying N tasks, each requiring time t to give evidence of their utility, with overhead H for switching during a fraction  $0 < \frac{Nt}{T} \leq 1$  of the allotted time T followed by exploiting the best one is  $U(F, N) = -H(N-1) + T(FE[u] + (1-F)E[u_*])$ . The difference from N=1 is  $\Delta U(N) = -H(N-1) + T((F-1)E[u] + (1-F)E[u_*]) = -H(N-1) + T(1-F)(E[u_*] - E[u])$ . There is some value of T where the last term becomes larger than the first term and it becomes rational to try different tasks. Since spending more time than necessary exploring wastes time best used for exploiting, F will be close to Nt/T. In the Gaussian case  $\Delta U(N) \approx -H(N-1) + T(1 - Nt/T)\sigma\sqrt{2 \ln(N)}$ , and the switch happens when  $T > N(H + t\sigma\sqrt{2 \ln(N)}) - H$ . Hence, as processes speed up (T increases) it becomes rational to flip between tasks, especially when overheads for task switching are low, the time to detect the utility well enough is fast, and the number of alternatives to consider remains bounded.



At least in principle, calculative rationality becomes perfect rationality if done infinitely fast (Russell 2016) (although as we will see this promise is somewhat illusory).

### *Competition*

In finance “winner take all” dynamics can occur in markets if one agent can react faster than others and hence gain a speed premium. The same may apply in evolution, where evolving faster than parasites or competitors is useful. This can serve as incentives for accelerating, even if it comes at a cost.

There would be an equilibrium when the cost of higher speed offsets the economic (or other) gains. At this point every actor runs as fast as is optimal. However, if the situation is a winner-take all situation the economic gains accrue mostly to the fastest actor, and it is rational to try to at least temporarily push into the “inefficient” speed frontier since slower competitors are pushed out of the market. Perceptions may also matter significantly for speed investment: there could be inefficiencies because agents overinvest in too fast systems.

In addition, Moore’s law and other expected speedups lead to design choices (like software bloat) that are suboptimal. Since agents can plan for a faster future they may risk overshooting the equilibrium by investing in speed that is expected to be optimal, but may actually be too ambitious.

### *Timing*

Sometimes it is more important to have a result or event at a *particular* time, and speed gives more control over the situation. For example, in bomb disposal one of the first steps is to vibrate the assembly – if there are any vibration detectors disarming it will be very hard, and if it detonates it will occur at a time of the disposal expert’s choosing. A bomb with a slow, unpredictable fuse will detonate at an unknown time, and this can be dangerous. Projects where one can ensure that at a predictable time one will know if it can succeed or that certain key steps will have been done will be less risky, reducing the risk premium.

### *Time budgets*

Everything that matters to us occurs within our time budget. These time limits may be set by human lifetimes, corporate lifetimes, the next budget period for an organisation, or the time until the situation changes too much.

An interesting special case is secrets, where the aim is to ensure that disclosure occurs *after* a fixed time occurs. Most secrets have a time horizon beyond which their importance has declined enough that it no longer matters if they are revealed. In cryptography this can be used to estimate required key lengths: given an assumed time horizon the key needs to be long enough that given optimistic predictions of future computing power it is not possible to brute-force the key. This has led to the concept of “transcomputational problems”, problems requiring more information processing than can be amassed even in principle. Many such estimates are based on common sense limits such as a computer covering the Earth, but the actual ultimate physical limits are far larger – yet there is rarely any challenge in constructing a sufficiently large problem to reach true limits.

Secrets also have half-lives due to random disclosure, leaks and espionage (Swire 2015). This is a more troublesome time limit since it is unknown (indeed unknowable): at best the secret-keeper can estimate the rough half-life and plan for that time period, as well as prepare for what to do if it is disclosed. Note that this makes cryptographic issues and computer speed far less important: a speedup of computation that makes the secret forceable earlier but still several half-lives into the future has no effect on the utility. In a leaky world perfect cryptography is irrelevant.

In software finding bugs typically requires monitoring/beta testing time proportional to the mean time between failure (MTBF). This can be speeded up by using more testers/users, but fast systems will also show internal bugs quickly. Interactive or environment-linked systems may however have long MTBF despite great internal speed: here the human or environment acts as the slow critical path in the computation and error-checking.

### *Time value*

The value of a remote event is increased (possibly from zero) by coming closer to the present. This includes becoming able to achieve something that previously was impossible to achieve within the given timeframe. For example, calculating weather forecasts in hours that previously would have taken years suddenly makes them useful.

Fast change also typically implies more uncertainty and faster discounting. If discount rates are pushed by accelerating computation but also affect human-level systems that do not change speed, there could be a problem in rational allocation. Long-term valuable projects such as building infrastructure or settling other planets are not only outside the fast discount horizon, but rapid changes in funding, technology, risk and specification makes planning challenging and reduces the probability that they will be fully implemented.

### *Early discovery*

Nick Bostrom penned “A little theory of problems” (Sandberg 2014) where among other things he noted that:

- “Discoveries” are acts that move the arrival of some information from a later point in time to an earlier point in time.
- The value of a discovery does not equal the value of the solution discovered.
- The value of a discovery equals the value of having the solution moved from the later time at it would otherwise have arrived to the time of the discovery.

In this account of the value of discovery depends on temporal ordering. In the long run most discoverable information will presumably be discovered, but problems with high and positive value *and* high elasticity (the solution can be found significantly sooner with one extra unit of effort) should be prioritized.

If computation is relevant for solving the problem improvements in speed means that it increases the elasticity, not just of the problem in question but for all computationally dependent problems. This can cause an overall re-prioritization: computational problems with a high value of early discovery would be favoured over other, equally weighty, but less elastic problems.

However, as long as computation is increasing in speed and power there may also be a value in waiting. As noted by Gottbrath et al. (1999), if a computation requires a given amount of computation to perform and the available computer power grows exponentially it can be rational to wait until the computer power has grown to start running the computation. This is relevant if the time the computation takes is significant compared to the improvement time of computers.

### *Cybernetic feedback shift*

When a regulatory or informative process becomes fast enough the dynamics changes profoundly.

We may call this a cybernetic feedback shift: when a controller is fast enough things become stable. This is well known from delay differential equations, the classic steam engine centrifugal governor, and trying to steer a boat: too slow reactions leads to oversteering where the system sways back and forth. In control theory delays are equivalent to a low sampling rate of the system. It may be possible to control the system, but the response will be sluggish since one has to use old data to stabilize. As control speeds up the responses become faster.

In many practical cases regulators try to track a parameter (e.g. thermostats, price signals). When the controller cannot follow a parameter, it tends to “fall off” with possibly catastrophic effects (Ashwin et al. 2012). For fast response rates this risk disappears. Conversely, if the *parameter* starts moving too fast because it has been speeded up, then stability may also be lost. If both controller and parameter speed up, then nothing changes.

In many cases the utility of a system goes up tremendously when a cybernetic feedback shift happens. Unfortunately the converse also occurs: when a system loses the ability to track parameters it can suddenly become worse than useless.

Communications matter for keeping organisations and empires together. An empire cannot function if the time from a province begins to rebel to the arrival of the military force sent from the capital is longer than the time the province needs to entrench/build army (a way around this is to distribute power to local governors for fast response). If the time it takes for the strategic level of the organisation to notice, understand, and react to a challenge is longer than the evolution time of the challenge, then it is unlikely to be able to deal with it. Hence real societies and organisations tend to have low-level, faster subsystems.

If local processing speeds up but communications does not this leads to asynchrony issues. But also, to stronger reasons to decentralise to meet new, fast challenges. Organisations that cannot do this will have a risk that “provinces” can rebel much faster and that subsystems are going to be held back – there may now exist local incentives to break loose.

## Risks and ethical issues due to fast computation

We have seen that there exist numerous reasons to pursue faster computation, especially in the service of performing particular tasks fast, timely, and early both for competitive and intrinsic reasons. Even if this is achieved flawlessly there are risks and ethically relevant issues with this process.

The most obvious issue of time compression is simply social change and disruption (as well as gains); the *locus classicus* is Toffler's *Future Shock* (1970) and the writings of Virilio; see also (Wajcman 2008).

Although clearly important, for reasons of brevity the focus here will be on the directly speed-related issues rather than the more subtle sociological and phenomenological challenges.

### *Loss of control*

The most obvious issue with fast computation or other operations is that they are fast relative to the (human) ability to control systems: there is a risk of things going haywire too fast.

When a system has a positive feedback loop, the strength of the feedback relative to friction and delays determines the speed a disturbance gets amplified. Engineering typically wants to keep friction in useful range: not too much to cause losses but enough to make system controllable. Faster computation and communication can make feedbacks in business, software (e.g. "Warhol worms" (Weaver 2001)) and other system more intense, and hence limit the timespan for taking control action. The "friction" – costs of action and reaction – is lowered, making fast actions possible.

Loss of control can happen due to several causes. There is the direct loss of control, where the steering agency lacks the ability to control, either because it cannot track the state of the system fast enough or process what should be done fast enough. There is the effect of emergence causing misbehaviour (systemic risk), where parts of a system function well but the whole exhibits unwanted behaviour. Here the trouble may come from the speed the emergent behavioural change occurs on, or that fast and dense internal interactions enable the change (Goldin & Mariathasan 2014). Finally, there is asynchrony where the parts cannot coordinate necessary joint activity.

Perhaps the most extreme example of an emergent loss of control is the "hard AI take-off scenario". In this scenario general artificial intelligence (AGI) is developed to have enough ability to perform many human-level tasks, including programming better AGI. A feedback loop ensues, where human input for improving AGI becomes less important than AGI input (which may be very scalable) and the total ability and power of the software grows rapidly, soon outstripping human ability to control. (Good 1966; Bostrom 2014) Whether such take-off can occur and how fast it could be remains conjectural at present, but it is an issue taken seriously as a long-term risk by some AI researchers (Hutter 2012; Müller & Bostrom 2016; Sotala 2017).

Bostrom distinguishes between fast, medium and slow take-offs in terms of how fast the transition occurs relative to human decision-timescales. The key issue is the differences in what reactions can be undertaken: in slow take-offs there is ample time for society to respond with considered actions, while in fast take-offs events move too quickly for human decisions to matter. In the intermediate case there may

not be time enough for deeply considered decisions, but various actions are possible – especially pre-planned “cached” actions that can be initiated quickly.

### *Speed inequality*

Speed differences can become unfair differences in economics or power, as well as contribute to risk.

In “fastest takes all” competitive situations being faster is more important than being good. This can favour not just excessive speedups and arms-races, but also ignoring quality and safety. For example, if several teams race to create the first transformative AI but safety work slows progress, then the Nash equilibrium tends to produce unsafe AI (and having public information on the progress of other teams increases the risk) (Armstrong, Bostrom & Shulman 2016). A less dramatic case is how Silicon Valley competition favours bringing a Minimum Viable Product to market fast and first rather than making it reliable; the result is often that security and privacy flaws become hard to fix later.

Old systems tend to be slower and would hence suffer in “fastest takes all” situations. Agents that cannot afford faster systems will hence tend to fall further behind. The speed requirements also serve as a barrier to entry.

Automated trading became possible in the 1990s when trading floors were replaced by matching engines. Gradually high-frequency trading emerged, as the second-long delays at the turn of the century declined to milliseconds in 2010, enabling trading of shares under 10 milliseconds. Quickly this became a dominant form of trading (Massa 2016) to a large degree because of improvements in liquidity and informativeness of quotes (Hendershott, Jones & Menkveld 2011). There are also more zero-sum benefits of speed such as obtaining a better position in the order book queues than competitors with the same information and similar strategies; the rewards for a 1 millisecond speed advantage have been claimed to be in the range of hundreds of millions to billions of dollars (Farmer & Scouras 2012). Human traders obviously cannot compete.

Also, the algorithms have shown sensitivity to disinformation/misinterpretation of news; oil prices jumped 2013 when a tweet recalling the Yom Kippur war 40 years later was sent by the Israeli military (Reuters 2013): since fast response is important double-checking signals is too slow, and once the market dynamics is set it becomes irrational to act just on the true information. This can contribute to instability, both in the large but also in the small in the form of ultrafast extreme events that are far outside human reaction times. There appears to exist a systemic transition when the number of agents is larger than the number of strategies and there is not enough time to process information (Johnson et al. 2013). Together, this both suggests unfairness caused by speed differences and risks from lack of control.

Speed inequalities matter in communication too: people need to be able to meaningfully respond to each other to have relations, and this includes being part of society. In an interaction it is the slowest participant that sets the overall speed. This may contribute to an incentive for fast systems to mostly talk to fast systems, and limit human contact. An extreme example is the social stratification predicted in a society of minds running at different speeds (Hanson 2016), but milder examples abound of side-lining slow responders in organisations and engineers minimizing requests to slow subsystems.

## *Decision speed*

Faster computation promises potentially near-instant decisions. However, these are subject to information limits: a decision can only happen when it is known with enough certainty that a triggering condition has occurred. Computation was sometimes in the past the slowing factor, but slow information arrival (due to sensors, communication speed, low bandwidth) is likely more relevant in a fast computation world. Acting early with less information is less certain and this will introduce risk. The real bound on decision speed may hence be acceptable uncertainty in a given situation.

“Faster computation” is typically measured relative to fixed rate of human activity. More things become possible inside our time budgets, but we cannot observe/control too fast activities directly. We can leave this to automation, but now we have delegation problem. Circuit-breakers for financial markets will stop pre-described events but may let through anything else, creating a false sense of security.

Many human systems have layers acting on different timescales, for example slow-changing constitutions underlying laws, policies, social norms and fashions. This ensures that observation and control can function. If breaks between the layers emerge, that means there is no meaningful control. This may have been a contributing factor to the financial crisis in that regulators did not understand the changing financial instruments and their implications.

Decision speed is also competing with time for deliberation. Drone pilots have the problem that faster systems make the human in the loop the slowest and most performance decreasing factor, while being morally responsible takes time (and the pilot will be blamed for low performance *and* morally questionable actions). This issue gets writ large in the case of nuclear missiles. It is about 30 minutes Moscow and Washington as the ICBM flies, just about time for a “red phone” call for negotiations to occur. But between Islamabad and Delhi it is 5 minutes: were an unauthorized or accidental launch to occur the time for internal and external deliberation is exceedingly short.

## *Loss of opportunity*

The irreversibility of time has great ethical importance. Choosing now and fast can remove opportunity for later choice, especially when actions are irreversible such as using up non-renewable resources or releasing information (Bostrom, Douglas & Sandberg 2016). This is an issue since we will likely have better information in the future and may hence evaluate the actions differently.

Yet waiting in a risky state can be worse than taking a risk since risks are likely to catch up with us. We may want to trade a “state risk” such as nuclear war being possible for a “transition risk” that removes the state risk at a price of a temporary but greater risk (a radical disarmament deal, inventing superintelligent AI to “solve geopolitics”) (Bostrom 2014, p. 233). State risks are reduced if we speed up macro-structural development: less time in risky periods. Transition risks are reduced by having more time for preparing. Problems where learning from experience dominate solutions benefit from slowing down and getting more time to have experience, while problems requiring forethought or insight may benefit less.

## *Identity over time*

A final issue is the accumulation of change. Complex adaptive systems interacting with the world will change their internal structure as a response: learning/forgetting information, restructuring itself, breaking or reproducing. This can change their fundamental identity in relevant ways. What constitutes the important aspects of identity depends on the system and observer, but the “wrong kind” of identity change must be avoided since it loses accrued value.

It is not the span of time that matters but the amount of change. Typically, there is a “change budget” that can soak up the modifications without losing identity. Software, people and organisations that instantly change often change identity in the wrong way, while the same transformation done gradually may be both identity preserving and acceptable to stakeholders. This is at the core of Toffler’s concept of future shock (Toffler 1970).

Faster information processing means a higher rate of accumulating change, straining the change budget. Rapid adaptation may be beneficial from a control perspective but risks reducing the change budget.

## **Societal impact and governance**

As we have seen, the great challenge outlined in this paper is growing gaps of speed between computation and the human (ethical gaps) and gaps of speed between computation and control (cybernetic gaps).

In terms of governance the risk is that this produces a policy vacuum (Moor 1985, 2005). There will be a growing number of situations where there are no policies, yet actions must be taken. The time needed to conceptualize the situation and deliberate it remains on the human timescale. The eternal refrain that technology is outpacing ethics represents not just an ethical gap due to speed but also a cybernetic gap due to lack of information: decisions cannot rationally happen faster than there is information to decide upon. The Collingridge dilemma (Collingridge 1982) is partially cybernetic .

Yet faster computation also strongly increases the power of the state and institutions through cybernetic feedback shifts and increased legibility (Scott 1998). Surveillance power has grown faster than Moore’s law since it scales with hardware, software, data availability *and* sensor ubiquity (while institutional oversight has hardly kept up at the same rate; there is a widening cybernetic gap here). If searching for a person can be done in real-time it is very different from an ongoing bureaucratic process.

Can governance be speeded up meaningfully, assuming human speed remains constant? (We will here ignore the possibilities of enhanced posthuman speed discussed in (Hanson 2016).)

Financial market circuit breaker rules are automated, speeded up governance solving the human-machine speed difference problem in a small domain. However, their utility depends on whether the system can detect the right issue. They need a decision parameter that is both necessary and sufficient for a break (and high quality). If it is not sufficient they produce false alarms. If it is not necessary, they might

miss things that matter. If it is not of high quality, it may measure the wrong thing. In the stock market measuring trading volume and value makes sense but may still miss subtle qualitative shifts (e.g. in correlations) predicting a systemic risk. This is basically a principal-agent problem, where the agent may be a thing designed for a purpose – but as Bostrom shows, the AI principal agent problem is doubly hard (Bostrom 2014).

Some parts of governance just consist of processing information or doing formal, well-defined decisions: these can in principle be speeded up. There is a drive to do this to save money, provide timely service, increase fairness etc. This will likely work best for routine, well-specified governance without elements of social or strategic intelligence. The risk is that we end up with algocracy, opaque decision-making with little legitimacy (Danaher 2016). Another risk is that since such routines will have apparent or real instrumental value they will be favoured over messy and slow routines requiring intelligence, producing a legible but far more limited governance system.

However, governance can also affect the speed of computation. While controlling the growth of technology in general is unlikely because of its broadness, epistemic unpredictability and utility, it is certainly possible to mandate speed limits (e.g. delays on stock markets to avoid high-frequency trading (Farmer & Skouras 2012)), mandating response time limits, or even mandating the use of certain technology (e.g. accessibility requirements for websites).

Another approach is the differential technology development principle (Bostrom 2014, pp. pp. 229–237): if potentially harmful technologies are developed more slowly than technologies reducing their risks, their benefits will become available with less risk. This can be achieved by focusing funding and research priorities on the harm reduction technologies. In the current context this might include promoting technical solutions to ethical and cybernetic gaps.

Controlling algorithms is hence not so much about banning practices as having an adaptive, global learning system that observes what is going on, remembers past states, maintains a set of values that may be updated as information arrives, and then changes incentives to promote these values – open to updating every part in this system.

## Conclusions

There is very much more computational speed to be had, and we will likely reach it. This will generally produce much more value – better productivity, better predictions, better control, more opportunities. But those desirable aims will also lead to control gaps, systemic risks, speed inequalities, overly fast or uncertain decisions. This will challenge governance strongly, risking both policy vacuums, a drive towards algocracy, and numerous principal agent problems in bridging ethical and cybernetic gaps. Yet strong enough governance can mandate speeds, and adaptive and distributed governance can update faster. It is possible to prioritize areas where speed issues are known to generate trouble for regulation and using differential technology development to stimulate foresightful and responsible technology development.



## Acknowledgments

This paper was inspired by Amnon Eden's questions. I would like to thank the participants in the workshop "From algorithmic states to algorithmic brains" at NUI Galway 22<sup>nd</sup>-23<sup>rd</sup> September 2016 and two anonymous referees for much feedback. I would also like to thank Örjan Ekeberg for first making the point about time as a non-renewable resource. This paper was supported by the ERC research project UnPrEDICT: Uncertainty and Precaution: Environmental Decisions Involving Catastrophic Threats, and partially written during my time at the Gothenburg Chair Programme for Advanced Studies 2017.

## References

- Amdahl, G. M. (1967). Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference* (pp. 483-485). ACM.
- Armstrong, S., Bostrom, N. and Shulman, C., 2016. Racing to the precipice: a model of artificial intelligence development. *AI & society*, 31(2), pp.201-206.
- Ashwin, P., Wieczorek, S., Vitolo, R. and Cox, P., 2012. Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system. *Phil. Trans. R. Soc. A*, 370(1962), pp.1166-1184.
- Bekenstein, J. D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2), 287.
- Bostrom, N., 2014. *Superintelligence: paths, dangers, strategies*. Oxford University Press.
- Bostrom, N., Douglas, T. and Sandberg, A., 2016. The Unilateralist's Curse and the Case for a Principle of Conformity. *Social epistemology*, 30(4), pp.350-371.
- Collingridge, D., 1982. The social control of technology.
- Cowen, T. (2011). *The great stagnation: How America ate all the low-hanging fruit of modern history, got sick, and will (eventually) feel better*. Dutton. Penguin.
- Ctistis, G., Yuce, E., Hartsuiker, A., Claudon, J., Bazin, M., Gérard, J.M. and Vos, W.L., 2011. Ultimate fast optical switching of a planar microcavity in the telecom wavelength range. *Applied physics letters*, 98(16), p.161114.
- Danaher, J., 2016. The threat of algocracy: Reality, resistance and accommodation. *Philosophy & Technology*, 29(3), pp.245-268.
- Denning, P.J. and Lewis, T.G., 2016. Exponential laws of computing growth. *Communications of the ACM*, 60(1), pp.54-65.

Drexler, K. E. (1992). *Nanosystems. Molecular machinery, manufacturing and computation*. J. Wiley & Sons

Eden, A. H., Steinhart, E., Pearce, D., & Moor, J. H. (eds.) (2012). *Singularity Hypotheses*. Springer, Berlin, Heidelberg.

Farmer, D. and Skouras, S., 2012. *Review of the benefits of a continuous market vs. randomised stop auctions and of alternative priority rules (policy options 7 and 12)*. Manuscript, Foresight, Government Office for Science, UK.

Grace, K., 2013. *Algorithmic progress in six domains*. Technical report, Machine Intelligence Research Institute. Lloyd, S., 2000. Ultimate physical limits to computation. *Nature*, 406(6799), p.1047.

Goldin, I., & Mariathasan, M. (2014). *The butterfly defect: How globalization creates systemic risks, and what to do about it*. Princeton University Press.

Good, I.J., 1966. Speculations concerning the first ultraintelligent machine. In *Advances in computers* (Vol. 6, pp. 31-88). Elsevier.

Gottbrath, C., Bailin, J., Meakin, C., Thompson, T. and Charfman, J.J., 1999. The Effects of Moore's Law and Slacking on Large Computations. *arXiv preprint astro-ph/9912202*.

Hanson, R. (2016). *The Age of Em: Work, Love, and Life when Robots Rule the Earth*. Oxford University Press.

Hendershott, T., Jones, C.M. and Menkveld, A.J., 2011. Does algorithmic trading improve liquidity?. *The Journal of Finance*, 66(1), pp.1-33.

Hutter, M. (2012). Can intelligence explode?. *Journal of Consciousness Studies*, 19(1-2), 143-166.

Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J. and Tivnan, B., 2013. Abrupt rise of new machine ecology beyond human response time. *Scientific reports*, 3, p.2627.

Kamigaichi, O., Saito, M., Doi, K., Matsumori, T., Tsukada, S.Y., Takeda, K., Shimoyama, T., Nakamura, K., Kiyomoto, M. and Watanabe, Y., 2009. Earthquake early warning in Japan: Warning the general public and future prospects. *Seismological Research Letters*, 80(5), pp.717-726.

Koh, H. and Magee, C.L., 2006. A functional approach for studying technological progress: Application to information technology. *Technological Forecasting and Social Change*, 73(9), pp.1061-1083.

Krugman, P., 2010. The theory of interstellar trade. *Economic Inquiry*, 48(4), pp.1119-1123.

Kurzweil, R. (2010). *The singularity is near*. Viking.

Lachmann, M., Newman, M. E., & Moore, C. (2004). The physical limits of communication or Why any sufficiently advanced technology is indistinguishable from noise. *American journal of physics*, 72(10), 1290-1293.

Mack, C. (2015). The multiple lives of Moore's law. *IEEE Spectrum*, 52(4), 31-31.

Margolus, N. and Levitin, L.B., 1998. The maximum speed of dynamical evolution. *Physica D: Nonlinear Phenomena*, 120(1-2), pp.188-195.

Markosian, N., (2016). "Time", *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.),

Massa, A. (2016). High-Speed Firms Now Oversee Almost All Stocks at NYSE Floor *Bloomberg* January 26, 2016

Mii, Y., Rishton, S., Taur, Y., Kern, D., Lii, T., Lee, K., Jenkins, K.A., Quinlan, D., Brown, T., Danner, D. and Sewell, F., 1994. Experimental high performance sub-0.1/ $\mu\text{m}$  channel nMOSFET's. *IEEE electron device letters*, 15(1), pp.28-30.

Moore, C., & Mertens, S. (2011). *The nature of computation*. OUP Oxford.

Moor, J.H., 1985. What is computer ethics?. *Metaphilosophy*, 16(4), pp.266-275.

Moor, J.H., 2005. Why we need better ethics for emerging technologies. *Ethics and information technology*, 7(3), pp.111-119.

Müller, V. C., & Bostrom, N. (2016). Future progress in artificial intelligence: A survey of expert opinion. In *Fundamental issues of artificial intelligence* (pp. 555-572). Springer, Cham.

Nagy, B., Farmer, J. D., Bui, Q. M., & Trancik, J. E. (2013). Statistical basis for predicting technological progress. *PloS one*, 8(2), e52669.

Nielsen, J., 1998. Nielsen's law of internet bandwidth. Online at <http://www.useit.com/alertbox/980405.html>

Pires, D.P., Cianciaruso, M., Céleri, L.C., Adesso, G. and Soares-Pinto, D.O., 2016. Generalized geometric quantum speed limits. *Physical Review X*, 6(2), p.021031.

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological review*, 106(4), 643.

Reuters (2013). Tweet recalling Yom Kippur war, 40 years on, jolts oil traders. October 10, 2013.

Russell, S. (2016). Rationality and intelligence: A brief update. In *Fundamental issues of artificial intelligence* (pp. 7-28). Springer, Cham.

Sandberg, A. (1999). The Physics of Information Processing Superobjects: Daily Life Among the Jupiter Brains. *Journal of Evolution and Technology*. 5(1), 1-34.

Sandberg, A. (2013). An Overview of Models of Technological Singularity. In *The Transhumanist Reader: Classical and Contemporary Essays on the Science, Technology, and Philosophy of the Human Future*, eds. More, M., Vita-More, N. pp 376-394.

Sandberg, A. 2014. What questions should we ask? Talk given at the second 2014 Oxford symposium on Big Data Science in Medicine, December 8 2014. *Online at* <http://www.aleph.se/presentations/What%20questions.pdf>

Schaller, R. R. (1997). Moore's law: past, present and future. *IEEE spectrum*, 34(6), 52-59.

Schwartz, B. (2004). *The paradox of choice: Why more is less*. New York: Harper Perennial.

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.

Sotala, K. (2017). How feasible is the rapid development of artificial superintelligence?. *Physica Scripta*, 92(11), 113001.

Swire, P., 2015. The Declining Half-life of Secrets and the Future of Signals Intelligence. *New America*.

Toffler, A. (1970). *Future shock*. Random House.

Wajcman, J. (2008). Life in the fast lane? Towards a sociology of technology and time. *The British journal of sociology*, 59(1), 59-77.

Waldrop, M. M. (2016). The chips are down for Moore's law. *Nature News*, 530(7589), 144.

Weaver, N., 2001. Warhol worms: The potential for very fast internet plagues. *Online at* <http://www.iwar.org.uk/comsec/resources/worms/warhol-worm.htm>