

IEICE **TRANSACTIONS**

on Communications

DOI:10.1587/transcom.2017ADI0003

Publicized:2018/02/22

This advance publication article will be replaced by the finalized version after proofreading.

A PUBLICATION OF THE COMMUNICATIONS SOCIETY



The Institute of Electronics, Information and Communication Engineers
Kikai-Shinko-Kaikan Bldg., 5-8, Shibakoen 3chome, Minato-ku, TOKYO, 105-0011 JAPAN

Autonomous Decentralised Systems and Global Social Systems

Colin G. Harrison[†], *non-member*

SUMMARY

As the capabilities and costs of Artificial Intelligence (AI) and of sensors (IoT) continue to improve, the concept of a "control system" can evolve beyond the operation of a discrete technical system based on numerical information and enter the realm of large-scale systems with both technical and social characteristics based on both numerical and unstructured information. This evolution has particular significance for applying the principles of Autonomous Decentralised Systems (ADS) [1]. This article considers the possible roles for ADS in complex technical and social systems extending up to global scales.

key words: autonomous decentralised system, ADS, complex systems, smart cities

1. Introduction

We are in a moment in history comparable to the period around 1988-1994 when we could anticipate the emergence of a global network - then known as The Information Superhighway - but we had little concept of the impact that this network - ultimately the public Internet - would have. Today IoT and AI, building on the pervasive global communications platform, are being prepared for wider application [2]. Engineering applications such as self-driving or autonomous vehicles (AV), heterogeneous and continental-scale utility systems, and more advanced robotics systems are in well-developed prototypes and their commercial deployment is imminent. Autonomous Decentralised Systems (ADS) [1] are poised to become even more widely diffused into the world and this article offers reflections on the implications of this trend.

ADS was originally conceived for spatially-distributed technical systems such as railways and utility networks in which it could not be assumed that a human operator would be physically or virtually present or where the failure of a human operator could have disastrous consequences. The intent of ADS was to permit such technical systems a degree of self-control, particularly to protect human lives or to ensure the stability of the system. However this permission was limited to purely technical systems. While human beings might be present in such circumstances, they did not form part of the system being controlled. Today it appears increasingly likely that ADS will extend to systems in which human beings are principal actors. This

evolution would involve ADS in the control of systems that are not merely complicated, but are complex.

The technologies driving the evolution of ADS include pervasive wireless networks of increasing bandwidth, intelligent instrumentation that provides ever more detailed and time-critical information on fixed and mobile technical systems as well as on human activities, at least in urban areas, and AI that can make rapid, intelligent decisions based on local and remote information. Thus independent technical systems, for example AVs, can be expected to be capable of assessment of their local environments in hard real-time, while remote systems, for example regional traffic management systems, will have detailed knowledge of all the independent systems and of related human activities.

We may expect that such independent technical systems will act autonomously to achieve their individual goals through very local decision-making, such as navigation and collision avoidance where response times of the order of milliseconds are important, and that regional or global regulatory systems will manage the collective behaviour of the technical systems to achieve overall societal goals. Thus we are no longer only concerned with achieving the safe and efficient operation of technical systems. We must consider also the roles that such technical systems play in complex social systems [2].

This article develops the distinctions between complicated technical systems such as ADS and complex social systems that include individual and collective human behaviours. It then considers the mutual adaptation of ADS and social systems and how that may influence other mechanisms such as legal systems that exist to regulate social behaviours. Finally it considers specific global-scale systems, notably the global financial system, global manufacturing, and global environmental systems.

2. Complicated and Complex Systems

Consider a large electrical utility network, one of the most complicated types of technical systems. Historically electricity utilities have balanced overall supply and demand in their networks through energy

[†] The author retired in 2013 from the IBM Corporation, Armonk, NY, USA as a Distinguished Engineer.

trading and among generators and distributors. This supply-side management is achieved through online markets that are cleared every fifteen minutes. That is, generating utilities offer to sell energy in the coming fifteen minutes at a certain price. Distributing utilities will accept such bids, starting at the lowest price, until they fulfill their anticipated needs for the coming fifteen minutes. This is called "clearing the market".

This method is well-established, but faces new challenges with the widespread deployment of both industrial-scale and micro-scale renewable generators - typically photovoltaic panels or wind generators - as well as new loads such as Plug-in Electric Vehicles (PEVs) that recharge their batteries from the electricity grid. At the residential level the control of these numerous, highly-distributed systems is effected through Smart Meters that act as agents and buy and sell energy on behalf of their consumers. This includes both supply-side and demand-side management.

Clearing the market now becomes far more complicated. Distributors may be required by law to take all energy offered by residential generators. PEV owners may be able to specify a time by which their vehicles must be sufficiently charged. PEVs constitute heavy loads on medium- and low-voltage distribution networks and they are often clustered in business-district car parks or in residential areas. Hence these positive and negative energy flows in the distribution networks must be balanced under very large numbers of commercial, legal, and technical constraints.

Suppose that in a hypothetical distribution network, there are on average four nodes between each consumer and each industrial-scale or micro-scale generator and that each node has four branches. Further suppose that low-voltage distribution networks aggregate 1,000 consumers. Hence there are some 10,000 groups as illustrated in Fig. 1.. So every fifteen minutes we have some $4^4 \times 10^4 =$ roughly 2 million routes over which energy flows must be balanced.

This is a Complicated System. It is massive in scale, deals with vast numbers of constraints, operates with timescales of minutes, and is critical to the lives of a large population. But it is an engineering system that is heavily instrumented and can be decomposed using system engineering principles into large numbers of less complicated sub-systems each of which is deterministic. This is typical of the challenges that IoT systems face. If the overall system were to fail, it could be restored. It is in fact coupled to other systems such as the weather, the extractive industries, and to the national economy, but these couplings are relatively weak at the operational time-scales.

Now consider on the other hand the global banking system. When a consumer receives a bill from a retailer,

the consumer instructs his or her bank to make a payment to the retailer's bank account. In general the consumer's bank is not the same as the retailer's bank. A given bank must thus manage large numbers of debit and credit instructions with many other banks. To simplify the completion of these transactions, the banks do not forward these for settlement one by one. Instead each bank holds money in credit on behalf of its principal correspondents. That is, the banks lend money to one another that is used to make local settlements of such transactions. Periodically, perhaps once a day, they make one settlement exchange of the net credit or debit. If each bank has such arrangements with, say, a hundred other banks around the world, then it is clear that a substantial amount of their assets are on loan to other banks and in return it holds credit from these banks.

Lending is a social process based on trust. The inter-bank networks depend on mutual trust and so the banks and their regulators pay close attention to their stability. Each bank is autonomous in the sense that it is free to include the credit it holds for its correspondents in the assets it employs for lending to its clients or indeed to other banks. If a bank fails, its correspondent banks may lose some or all of the money loaned. There is a risk that one such failure may propagate through the network and lead to domino failures of other banks. (A similar knock-on failure mechanism can exist in large-scale utility networks.)

In general the bigger the bank, the bigger the impact of its failure would be on the banking network. This leads to the concept of being "too big to fail". That is, above some scale, national or international regulators would prevent the failure of a very large bank by injecting credit. Moreover, banks have the right to withdraw the money that they lend to other banks. So a loss of trust or confidence in a given bank, perhaps arising from that bank's bad loans, may cause its correspondent banks to exercise or to threaten to exercise this right, putting the bank under even greater pressure.

So we can visualise the global banking system as a network in which banks are the nodes and the edges represent inter-bank lending relationships as illustrated in Fig. 2. Associated with each edge is the lending party's assessment of the recipient bank's credit-worthiness, which is the ability of the borrowing bank to manage that loan [4]. This inter-bank network is critical to the operation and stability of the global financial system. Network analysis tools are applied to assess its overall resilience to potential bank failures.

Prior to the 2008 financial crisis it was believed by economists, financial analysts, and regulators that making this inter-bank network as highly connected as possible would ensure stability. That is, the more

highly connected any given bank is, the more likely it would be able to survive major losses in its own loan portfolio, because it has more correspondents from which to draw credit. This turns out to be wrong. As a member of the European Central Bank's advisory board expressed it: "Banks are not too big to fail, they are too connected to fail." That is, the failure of a bank in a highly-connected network may pose a strong risk of propagating widely and risks taking down the entire network. As correspondents lose confidence in a bank with apparent bad loans, they are apt to withdraw their own loans to the bank, thus increasing the likelihood that it will fail. But its failure and the resulting losses to other banks may cause one or more them to fail and so the high degree of connectivity causes even greater destruction.

This scenario is what almost happened in 2008 as a result of sub-prime lending in the USA and other parts of the world. The US regulators allowed one sizeable bank, Bear Sterns, to fail, but, as they realized the global risk, they then injected massive amounts of credit to prevent further failures.

The global financial system is not merely a Complicated System, it is a Complex Systems. Its instability results in large part from human judgment. It is not deterministic and not decomposable. The inter-bank network is a global Complex System, unlike the utility network which is merely a localised Complicated System.

The characteristics of Complex Systems include:

- Non-deterministic behaviour. The system does not respond reproducibly to a given stimulus.
- Non-linear behaviour. The Butterfly in China effect.
- Multiple factors influence behaviour. The system may contain immense numbers of interactions, although the number of significant interactions in a specific time and place will be small.
- Behaviours at multiple spatial and temporal scales. Analysis and interventions must consider at least one level above and below the scale of immediate interest.
- Initial conditions cannot be controlled. The system emerges from some previous state and exists because it exists.
- High degree of resilience. The system exists because it has evolved to survive (most) events.
- Tipping points. In extreme circumstances, the system may make a rapid transition to a different state or set of behaviours. It may be difficult or impossible to restore it to the

previous state. For example, species extinction is an irreversible process.

- No boundaries. The system cannot be isolated from multiple, global interactions. This makes it difficult, if not impossible, to apply scientific methods in which only a single factor is allowed to vary. Some systems may even extend beyond the planet; for example influences from the moon or the sun on climate and cycles of plant and animal life.

There are many Global Complex Systems including weather, climate and natural and social ecological systems such as habitats. They are difficult if not impossible to establish *ab initio*, but rather emerge and survive if they prove to be viable. If they are viable, they are likely to be strongly resilient to external changes (this is a circular argument). They will respond to such changes up to some level of stress in a quasi-linear and reversible manner, but beyond that level may flip to a different state from which it may not be possible to return to the initial state. A familiar example is the way in which traffic speed declines monotonically as traffic density increases and then falls to zero. If this results in gridlock, it becomes very difficult to restart traffic flows.

3. ADS in Social Systems

As the power of IoT and AI increase dramatically, their applications begin to impinge on social systems. This is not a new phenomenon. Every significant technological advance creates new dependencies in the societies that embrace it. Examples include the Haber-Bosch Process [5], which dramatically increased crop yields, leading to massive population growth, the Bessemer Process [6], which enabled the industrial-scale production of steel, resulting in railways, bridges, and multi-story buildings, or the development of Alternating Current systems by de Ferranti and Tesla, among many others [7], that launched the electrical utilities as we know them today, and not least the Internet, which has many parents, and has transformed many aspects of developed societies. Such innovations change the societies that adopt them in irreversible ways. Modern civilisation is riding many such tigers from which we are cannot off.

While it is true that these technologies and the industries they have fostered are connected to Complex Systems such as economies and natural and social environments, they are in themselves merely Complicated Systems and our civilisation has ever-increasing abilities for managing Complicated Systems. What is challenging is that we now see IoT and AI becoming embedded in Complex Systems such as road traffic and global banking. This has implications both for the engineers who are designing and operating such systems and for

policymakers who must plan for their introduction into society and for their regulation. Such complex Social ADS may have emergent properties for which neither engineers nor policymakers nor citizens are prepared.

Consider for example Autonomous Vehicles (AVs) and their relationships to complex social systems. Let us accept for the moment the claims made by the developers of these vehicles that they can match or surpass human drivers in terms of navigation, of the detection and avoidance of dangerous situations, and of compliance with the rules of the road. In other words that they can emulate or exceed human drivers in the technical aspects of driving¹. But human driving is not merely about regimented road manners. Urban driving is also a collective activity guided and constrained by local *mores*.

Readers will be aware that driving in New York is different from driving in San Francisco or Mumbai or Rome in matters such as entering a lane, making a left turn, responding to cyclists or pedestrians, and many others. Further, these are not the behaviours of individual drivers, but rather collective behaviours somewhat like the manoeuvres of a flock of birds. These behaviours may vary by neighbourhood or at individual road junctions. They may also vary by time of day, day of week or season, and by weather condition. They are individual actions in the collective or social system of road users.

Such social systems may be highly resilient, able to adapt and to return to "normal" when perturbations disappear. They are also highly effective at solving the complex problem of how some millions of people get to work or back to home every day with maximum (if not high) satisfaction. In terms of economic theory, we may say that each citizen attempts to maximise his or her own satisfaction with the city's many social systems. It is for this reason that these systems are highly resilient: to the extent that the actors act rationally, they have maximised their individual satisfactions.

AV technologies [8] view each vehicle as an independent entity whose only responsibility to others is to adhere to the highway code and to avoid accidents. The technologies envision fleets of purely technical systems proceeding down largely featureless streets, whereas the reality today is that driving is a collective exercise in which drivers compete for road space through learned patterns of behaviour, but also collaborate to a certain degree, reflecting the driving culture of the city. To what extent do the developers of AV technology understand how urban communities, which are certainly Complex Systems, work?

¹ The author does have reservations about the vast number of other use cases that have not been considered.

In addition to such local culture, social behaviours are also constrained by laws and regulations. Formal legal systems in settled communities go back several thousand years and were no doubt preceded by many thousands of years of informal *mores*. A legal system is a centralised control system that seeks to constrain individual behaviour in order to produce certain social benefits, for example safe roads, or more broadly, a safe and stable society. It does this by defining the boundaries of behaviour and then imposing penalties on individuals or organisations that transgress these boundaries. Although the definition of these controls is centralised at some level of government, its operation is distributed throughout society, for example by police patrols, and these patrols have a degree of autonomy or judgment. We will call this Social ADS.

It could be argued that such laws and regulations, which apply without exception to all citizens and organisations, reduce the society's capabilities by excluding exceptions that would probably not reduce the intended social benefit of, for example, safe roads, but could improve some other social benefit, for example, delivering hot food. Many societies tolerate violations of speed limits in driving, although the degree of tolerance varies widely. In part this is a practical matter - societies cannot (yet) afford to police every citizen's behaviour at this granularity - and in part an acceptance that citizens' judgments of "safe" behaviour - conditioned by *mores* - can be an adequate alternative.

The rise of interest in AI and IoT leads to speculation about their application in the operation of legal systems, although science fiction writers got there first in works such as "Minority Report" [9]. In the context of Social ADS, we might confer on individual systems, for example AVs, a certain, momentary discretion. For example, if there were no other vehicles within 100 m of an AV, it might be permitted to exceed the speed limit by, say, 10 km per hour.

However as part of a complex social system, we would also need to consider not only the impact of a single individual, but also the overall impact of permitting such transgressions. For example, the impact of a wave of such speeding vehicles arriving at an area of dense traffic. In fact AIs and IoT should eventually be able to deal with far more complex situations in which the driving behaviour of entire cohorts of vehicles, each of which has a stated destination and time constraint, is managed as a kind of massive board game with each vehicle being instructed on speed, lanes, exits, and so forth.

Such a legal approach has been mooted by Casey and Niblett [10, 11] who have suggested that, instead of defining desirable behaviour via laws and regulations that can be applied *post hoc* to assess behaviours for violations, an individual or a device operated by an

individual, notably an AV, would be given particular, real-time guidance on actions that would be considered legal under the present circumstances. They call these advisories "Micro-Directives".

Laws would still be enacted with the aim of achieving certain social goals, but rather than a single "one size fits all" approach, they would define multiple possible forms of compliance - the Micro-Directives - that would be based a wide range of scenarios. In a specific situation, the governance or control system would assess the context of an individual's situation and provide a Micro-Directive that will maintain compliance in the given scenario. Depending on the circumstances this may expand or reduce the operator's or device's freedom of action while maintaining the goal of a safe and stable society.

We might consider the existing legal system as an open-loop control system in which the output level of the system, for example an engine, is clamped well below its peak capability to ensure stability in a system where the operator cannot see the demands being placed on it. This could be, for example, a water pumping system where the operator cannot observe the instantaneous demand for water and must simply run the pump at a safe level that will on average meet the consumers' needs. The Micro-Directives approach introduces the ability to observe that demand and to allow the pump's control system to respond to peaks and troughs, while locally operating the pump safely and efficiently.

A scenario that is currently debated in the ethics of AVs concerns the problem of what such a vehicle should do if a pedestrian unexpectedly steps into the roadway. The vehicle is posited to have the option of quickly moving into the opposite lane, but it would then collide with an oncoming vehicle that is known to contain no human being. Since the centralised traffic management system is fully informed about the situation of every AV, it can assess this problem and issue a Micro-Directive that authorises or instructs the vehicle to change lanes and unavoidably collide with the on-coming passenger-less vehicle.

4. ADS in Global Social Systems

We may consider this Social ADS approach to laws and regulations as a consequence of Coase's Law [12]. Ronald Coase was one of the first economists to try to understand enterprises as complex systems in the early 20th century. At that time, large organisations, for example corporate enterprises such as the Ford Motor Company, would often be completely vertically integrated. That is, they would own and operate all the resources to manufacture their products. In Ford's case this extended as far as owning rubber plantations in order to produce tyres. In the computer industry such vertical

integration continued until *circa* 1980. Coase asked the question: What criterion determines when a corporation should outsource some product or service to a third party? His conclusion was that it is the relative costs of internal versus external transactions. If a company needs to produce a large quantity of tyres then it may be preferable to pay to own and operate a rubber plantation rather than to deal with the costs, latencies, and uncertainties of global rubber markets.

The advent of the Internet and IT support for global business to business (B2B) commerce in the 1990s dramatically changed the balance of these costs and lead, in part, to the rise of the globalised economy. It was speculated that large enterprises would be deconstructed into networks of specialised suppliers assembled into global supply chains. This expectation has been fulfilled only to a degree, as there are still advantages to being a large enterprise, notably in the ability to quickly raise capital for new products or markets.

Hence in the globalised economy, big industries often take the form of global supply chains. These are networks of many independent companies whose products and services are composed together by the manufacturers or producers of end-products and services. Each of these companies is autonomous in how it buys materials or services from its upstream partners, adds its own value to these, and sells its own products and services to its downstream partners. Such global supply chains have no centralised management system, but through applying ADS principles each aims to respond collectively to market demands.

Each participant in such a supply-chain must carefully sense down-stream demand and adjust its own upstream demand accordingly so that overall there is a smooth flow of products and services. It is somewhat like trying to manage the smooth flow of AVs in a city. Large companies, especially those at the end-product part of a supply chain will maintain situation centres that monitor the many possible events around the world that could disrupt such flows. These may be social - holidays, elections, political instabilities, strikes, financial crises - or environmental - storms, natural and man-made disasters, crop-failures - and their potential impacts on the specific supply chain are assessed and mitigating measures taken.

Most supply-chains incorporate alternate suppliers to ensure resilience, although major disasters, such as the 2011 tsunami in East Japan, may reveal dependencies on unique suppliers of a product whose unavailability can shutdown the entire supply-chain. Supply-chains may also manifest non-linear behaviour, such as "whiplash" in which a small perturbation, such as an increased demand for the end-product, propagates back up the chain and becomes amplified into a major demand for raw

materials.

5. Social ADS

Through the legal system, governments at various levels act in some ways as centralised, societal control systems with the aim of producing the "service" we call a safe and stable society. However, as the Micro-Directives example above illustrates, it is becoming possible to introduce higher degrees of flexibility and adaptation in how government achieves such goals. As noted above with respect to speed limits, legal systems already permit a degree of flexibility. However this flexibility is often *post hoc* in the form of the court system, which must assess, with considerable expense and delay, whether an accused person in a specific situation has truly violated the intent of a law or regulation.

Overall such flexibility is still currently more limited than it might be by the need for margins of safety due to the inability to observe situations in real-time. Societies have varying approaches to such flexibility. Some, for example the United States, prefer complicated laws that contain many exceptions. Whereas others, for example Switzerland, may prefer simpler laws with few or no exceptions.

Such Social ADS approaches could lead us to the *a priori* granting of flexibility that is wider than before across defined ranges of scenarios. These approaches would also require a changed relationship between citizens and government with a greater emphasis on citizens understanding the intent of the law, its social purpose, rather than understanding exactly where the legal boundaries lie. This change may be challenging for some to comprehend and certainly contentious, since a social purpose is a less precise concept than, say, a speed limit.

Such approaches may eventually extend from national laws to international laws. The dramatic expansion of the human population during the 20th century and the resulting expansion of agriculture and industries have also created global impacts on complex natural systems such as climate and the oceans. While humans do not (yet) directly participate in such complex natural systems, our activities have reached sufficient scale to produce strong indirect deleterious effects. With few exceptions, all the countries in the world have agreed to take steps under the Paris Agreement [13] with the goal of reversing these effects.

At present these steps are focused on reducing the total emissions of Green House Gasses (GHG), notably carbon dioxide, but also methane and certain industrial gases, and with reducing the destruction of forests, which are important for absorbing and sequestering carbon dioxide. While the current focus on carbon dioxide is appropriate, the global activities of humankind seem

likely to affect climate in other ways. Thus forests are important not only for sequestering carbon dioxide, but also for their beneficial effects on precipitation and on how rain is absorbed into the ground. The discharge of industrial and human waste into ecosystems at current scales is leading to phenomena such as the Great Pacific Garbage Patch [14], the extensive pollution of oceans with plastic materials, and to harmful algal blooms [15] in the Great Lakes of North America, the Persian Gulf and other regions.

At present these global phenomena are poorly understood as complex systems. Given the circulation of the atmosphere, in the short term it makes little difference where reductions in GHG emissions are made and each party to the Paris Agreement may proceed independently. But if the current trends can be brought under control and reversed and if the atmospheric system becomes more deeply understood, it is possible that in the coming century a global control system will emerge for managing allowable emissions in different regions. This may then resemble Social ADS as described above by introducing greater flexibility into local or national systems at a global scale. In concept this is distinct from various proposals mooted in recent years for climate engineering or geo-engineering [16], in which humankind would attempt to intervene directly in the climate system, although by the next century these may have coalesced.

6. Conclusion

The above discussion suggests a world in which the principles of ADS are applied not only to make technical systems autonomic and decentralised, but also to make some social systems autonomic and decentralised. This possibility emerges from the pervasive application of the Internet communications platform to collect and interpret both numerical and unstructured information. If this trend were pursued, it could change the structure of developed societies through providing greater visibility and understanding of context and hence giving greater autonomy and independence governance at the local and global levels.

However, this is by no means a certainty. The same technologies could equally well be applied to highly centralised control of dependent technical and social systems. We have already seen that the anticipated decomposition of large enterprises into many smaller enterprises has not occurred. Indeed the power of platforms such as Amazon, Facebook, Google, and Uber are powerful examples of this opposite possibility.

This balance between centralised and decentralised governance is a continuing theme throughout human history. Generally large communities have adopted some form of hierarchical governance. In part this tendency reflects the limited numbers of members who

were sufficiently educated and experienced to be leaders. This has changed somewhat with the availability of public education and professional training. The historic roles of monarchies have also been replaced by various forms of democracy.

Still this balance has moved backwards and forwards many times across many countries since the 18th century and this should be expected to continue. In the early 21st century, Information Technology is often accused of facilitating highly centralised government based on pervasive surveillance. The concepts of Social ADS show that it can also facilitate highly decentralised, autonomous government. But ultimately societies will need to decide on the balance between these two possibilities.

Figures




Figure 1: A hypothetical electricity distribution network showing generators (G) and consumers (C)..




Figure 2: A highly simplified illustration of the global inter-banking network at the time of the 2008 crisis.

Acknowledgment

The author gratefully acknowledges many stimulating discussions with Prof. Kinji Mori, Dr. Ralph Dum, and colleagues in the Global Systems Science and Policy community.

References

[1] Mori , Kinji (1993), *Autonomous decentralized systems: concept, data field architecture and future trends*, Proc. IEEE International Symposium on Autonomous Decentralized Systems (ISADS), 1993, pp. 28–34

[2] Harrison, Colin G., *Augmented Intelligence and Society*, in *Augmented Intelligence*, editor Daniel Araya, to be published by Peter Lang, 2018

[3] Harrison, Colin G., *Autonomous Decentralised Systems and Society*, Proc. IEEE 12th International Symposium on Autonomous and Decentralised Systems (ISADS), Taichung, Taiwan, March 2015

[4] Leduc, Matt V., Poledna, Sebastian, and Turner, Stefan, *Systemic risk management in financial networks with credit-default swaps*, (2017) , DOI: 10.21314/JNTF2017.034. See also https://en.wikipedia.org/wiki/Too_connected_to_fail

[5] https://en.wikipedia.org/wiki/Haber_process

[6] https://en.wikipedia.org/wiki/Bessemer_process

[7] https://en.wikipedia.org/wiki/Alternating_current

[8] https://en.wikipedia.org/wiki/Autonomous_car

[9] Dick, Philip K. (1991), *The Minority Report*, Citadel Twilight ISBN 0-8065-1276-8

[10] Casey, Anthony; Niblett, Anthony (2016), *Self-Driving Laws*, University of Toronto Law Journal, vol. 66, no. 4, <http://dx.doi.org/10.3138/UTLJ.4006>

[11] Casey, Anthony; Niblett, Anthony (2015), *The Death of Rules and Standards*, Available at SSRN 2693826.

[12] https://en.wikipedia.org/wiki/Ronald_Coase

[13] The Paris Agreement, http://unfccc.int/paris_agreement/items/9485.php

[14] https://en.wikipedia.org/wiki/North_Pacific_Gyre

[15] https://en.wikipedia.org/wiki/Great_Pacific_garbage_patch

[16] https://en.wikipedia.org/wiki/Algal_bloom

[17] https://en.wikipedia.org/wiki/Climate_engineering



Colin **Harrison** retired in 2013 as an IBM Distinguished Engineer. He was most recently the inventor of IBM's Smarter Cities technical programme. He was previously Director of Strategic Innovation in IBM Europe and Director of Global Services Research. He is an IBM Master Inventor and a Member of the IBM Academy of Technology. He writes and speaks on smart cities, the societal impacts of Information Technology, and global systems science, and searches for an underlying science of cities.