International Transport Forum

CPB
Corporate Partnership Board

Blockchain and Beyond: Encoding 21st Century Transport

Corporate Partnership Board Report

OECD

Blockchain and Beyond: Encoding 21st Century Transport

Corporate Partnership Board Report

About the International Transport Forum

The International Transport Forum at the OECD is an intergovernmental organisation with 59 member countries. It acts as a think tank for transport policy and organises the Annual Summit of transport ministers. ITF is the only global body that covers all transport modes. It is administratively integrated with the OECD, yet politically autonomous.

ITF works for transport policies that improve peoples' lives. Our mission is to foster a deeper understanding of the role of transport in economic growth, environmental sustainability and social inclusion and to raise the public profile of transport policy.

ITF organises global dialogue for better transport. We act as a platform for discussion and pre-negotiation of policy issues across all transport modes. We analyse trends, share knowledge and promote exchange among transport decision makers and civil society. ITF's Annual Summit is the world's largest gathering of transport ministers and the leading global platform for dialogue on transport policy.

Our member countries are: Albania, Argentina, Armenia, Australia, Austria, Azerbaijan, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Canada, Chile, China (People's Republic of), Croatia, Czech Republic, Denmark, Estonia, Finland, France, Former Yugoslav Republic of Macedonia, Georgia, Germany, Greece, Hungary, Iceland, India, Ireland, Israel, Italy, Japan, Kazakhstan, Korea, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Mexico, Republic of Moldova, Montenegro, Morocco, Netherlands, New Zealand, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, Ukraine, United Arab Emirates, United Kingdom and United States.

www.itf-oecd.org

Foreword

The work for this report was carried out in the context of a project initiated and funded by the International Transport Forum's Corporate Partnership Board (CPB). CPB projects are designed to enrich policy discussion with a business perspective. They are launched in areas where CPB member companies identify an emerging issue in transport policy or an innovation challenge to the transport system. Led by the ITF, work is carried out in a collaborative fashion in working groups consisting of CPB member companies, external experts and ITF staff.

Many thanks to the members of the Corporate Partnership Board companies involved in this work: Abertis, Alstom, Anheuser-Busch InBev, Ford Motor Company, Kapsch TrafficCom, Kapsch TrafficCom, NXP, PTV Group, Renault-Nissan Alliance, SAS, Siemens, SNCF, Toyota Motor Corporation, Uber, Volvo Car Corporation.

The report draws conclusions from the ITF CPB Workshop "Blockchain and beyond: Encoding 21st Century Transport" held 8-9 November, 2017 and hosted by Professor Sandy Pentland at MIT Media Lab. The meeting was chaired by Yves-Alexandre de Montjoye. Participants of the workshop included:

Chris Ballinger, Toyota Research Institute	Ruth Miller, Apple
Brandon M. Belford, Apple	Sho Nabeshima, Toyota Motor Company
Lily Bui, Massachusetts Institute of Technology	Neha Narula, Massachusetts Institute of Technology
Maguelonne Chandesris, SNCF	Sam Penfield, SAS
Denis Darmouni, Renault Nissan	Sandy Pentland, Massachusetts Institute of Technology
Marisa Anne DeAngelis, Massachusetts Institute of Technology	Christian Reimsbach-Kounatze, OECD
Benjamin de la Pena, Seattle Department of Transportation	Michael Replogle, New York City Department of Transportation
Yves-Alexandre de Montjoye, Imperial College London	Marc Ribo Pedragosa, Abertis
Lee Ann Dietz, SAS	Shaleen Srivastava, PTV-Group
Justin Erlich, Uber	Karen Vancluysen, Polis
Jenn Halen, Harvard University	Ellis Verosub, Apple
Paolo Humanes, PTV Group	Kevin Webb, SharedStreets, Open Transport Foundation
Tuomas Kaivol, Finnish Ministry of Transport and Communications	Guenther Wildmann, Kapsch TraffiComm
Holly Krambeck, World Bank - Open Transport Foundation	Sarah Williams, Massachusetts Institute of Technology
Paul Leghart, emovis	Philippe Crist, International Transport Forum
Ian MacBeth, Transport for London	Katja Schechtner, International Transport Forum/ Massachusetts Institute of Technology

Special thanks go to Ande Monier, Cécilia Paymon and Edwina Collins from the ITF and Marisa DeAngelis from MIT for notes, support and editing. The principal author of this report is Philippe Crist with contributions from Katja Schechtner. The project was managed by Philippe Crist and Katja Schechtner of the International Transport Forum. Coordination of CPB activities was carried out by Sharon Masterson.

Table of contents

Figures

Executive summary

What we did

Digital technology continues to reshape the transport industry. Recently, much discussion has focussed on blockchain and other distributed ledger technologies (DLTs). This report investigates the potential for DLTs to support broader coordination of seamless urban mobility services and the delivery of Mobility as a Service (MaaS) in urban settings. Like other economic sectors, transport could be profoundly transformed by blockchain, and other novel DLTs that allow decentralised applications to run in peer-to-peer networks.

These technologies allow agents to enter into direct relationships with each other according to a commonly agreed set of rules and a high degree of trust without having to go through a central authority. Combined with a common language and syntax for the "internet of mobility" and new means of deriving insight from previously siloed data, these applications may help redefine how people access, pay for and use transport in their everyday lives.

This report builds on an expert workshop at the MIT Media Lab in November, 2017 and further expert inputs and desk research carried out by the ITF Secretariat. It serves to frame the principal policy considerations relating to the application of distributed ledger technologies such as blockchain to an evolving urban mobility ecosystem.

What we found

Urban mobility today is a siloed world of separate and independently regulated services. The application of distributed ledger technologies, such as blockchain, to urban mobility may lead to a future more aligned with other "as-a-service" models where actors engage directly with each other based on commonly agreed protocols.

Even actors that have disrupted traditional transport (and other sectors) in recent years may in turn find their business models under pressure as citizens gain direct control to build their own trip experiences. These changes will also challenge public authorities. They must keep abreast of developments in data science and DLTs to adapt current regulations where they hinder beneficial outcomes. They must also explore new regulatory responses where these are necessary to deliver the outcomes the public wants.

The deployment of DLTs is still very much in its infancy, especially in support of Mobility as a Service (MaaS). Initial use cases for these technologies will not necessarily be those that get adopted at scale later. It is yet unclear if, how and to what extent DLTs will become integrated into economic sectors including transport and, ultimately, into daily life. Uptake hinges on whether or not decentralised ledgers such as blockchains can deliver better value than traditional ledger and transaction frameworks in use today. It will also depend on whether they can enable new, value-adding applications that are not yet possible with existing technologies, and on how far the regulatory environment will support this innovation.

What we recommend

Public authorities must prepare for a much more networked and meshed world

In an ever-more dynamic urban mobility ecosystem, citizens' choices are expanding and changing rapidly. Public transport operators, car manufacturers and taxi companies are facing increasing pressure to innovate to attract and to retain users. Increasingly they must both compete and co-operate with each other as well as with market entrants with new business models. Traditional regulatory approaches that focus on transport operators and modes in isolation are increasingly out of step with what recent market offers and how many people make travel decisions. Public authorities must adapt their regulatory framework to this

emerging and interconnected "mesh-y" urban mobility ecosystem. Legislation has to set the framework for interoperable MaaS but technical details must be addressed through standardisation bodies. The process for setting these standards must be inclusive, transparent and technically thorough.

Take into account changes in data science and technology when developing Mobility as a Service

The concept of Mobility as a Service (MaaS) encompasses the integration of various forms of transport services into a single mobility service, increasingly also accessible on demand. It offers people seamless digital access to different transport services, many of them shared. MaaS applications leverage database, identity management, data access and transmission protocols that are all simultaneously evolving. Current platform-based models for MaaS may not deliver on the promise of an open urban mobility ecosystem. In an increasingly networked and meshed world, a very diverse set of stakeholders must trust each other and underlying business processes. New developments in data protocols, structure and data science may help establish this trust in an open and platformless world. Governments should bolster their capacity to identify, understand and monitor these developments and support the implementation of the most promising of these.

Look beyond initial cryptocurrency applications of distributed ledger technologies

Much of the discussion around DLTs centres on their role to underpin cryptocurrencies. In the context of transport, this focus is misplaced. The value of DLTs for transport is how they ensure transparency, traceability, trust and distributed revenue sharing and governance. DLTs create new business and regulatory opportunities in a number of ways. They underpin robust identity and rights management. They create an immutable, distributed and openly verifiable record of past transactions. They enhance data privacy and access. And they improve cyber-security. DLTs also foster innovation and efficiency via automated business or regulatory processes by self-executing "smart contracts". Public authorities should identify the opportunities for better regulation and service delivery created.

Governments should help deploy the building blocks that enable wider uptake of distributed ledgers

As a technology and as a foundation for new business and regulatory processes, DLTs are still in their early days. It is hard to predict how this will evolve. This makes it difficult for public authorities to assess what role DLTs will play in MaaS - and what role public authorities will play in deploying DLTs, if any. Rather than supporting broad-scale deployment of existing DLTs, public authorities could ensure that the necessary building blocks are in place for future DLTs. These could include harmonised identifiers, a shared and common data syntax in support of the internet of mobility and a regulatory framework that anticipates future DLT developments. These standardised building blocks are effectively public goods and governments should use their convening power to bring developers together to establish them early.

Apply blockchain technology now for slow and (relatively) small transport use cases; anticipate next generation distributed ledger technologies for "big and fast" applications to be deployed later

Current blockchain applications are limited because they fail to scale and are relatively slow. Nonetheless, they are still suited to some of the tasks in delivering MaaS, those that are not sensitive to limitations in capacity to handle data, volume, or speed of processing. These include identity management, licensing and registration and asset tracking. These use cases can serve as a test bed that will allow stakeholders to become familiar with DLT-supported MaaS applications. MaaS tasks that require more real-time logging and high-volume data processing will require new DLT models purpose-built for speed and "Internet of things" applications. Technologies should already be tested even if their large-scale uptake for MaaS may not be immediate.

Governments should develop algorithmic code-based regulation to accompany the uptake of distributed ledger technologies

Transport activity is increasingly influenced by an underlying web of code-based algorithms and protocols. The deployment of DLTs in transport will only amplify that trend. At a minimum, public policy should understand how algorithms are affecting transport. But governments must also explore ways to move away from sole reliance on analogue, paper-based regulation that is crafted in human language. Instead, they will need to move towards frameworks that integrate technical code and algorithmic logic into the regulatory process – e.g. RegTech. This will require governments to enhance their internal capacity to understand and regulate in this domain, including through machine-to-machine communication.

Introduction

Decentralised applications running in peer-to-peer networks built on distributed ledger, blockchain and other novel data protocols are starting to profoundly disrupt established economic sectors (e.g. finance, healthcare, provenance authentication, commerce). These applications allow agents to enter into direct relationships with each other according to a commonly agreed set of rules and a high degree of trust without having to go through a central authority. Combined with a common language and syntax for the Internet of mobility" and new means of deriving insight from previously siloed data, these applications may help redefine how people access, pay for and use transport in their everyday lives.

The pace of change is such that even those actors that have themselves disrupted traditional transport (and other sectors) in recent years may quickly find themselves disrupted in turn as citizens gain direct control to mediate and shape their own mobility. These changes will also challenge public authorities who must keep abreast of such developments to adapt current regulations where they hinder the delivery of beneficial outcomes and explore new regulatory responses where they are warranted by the public good.

Crucially, these developments are still very much in their infancy – initial use cases for these technologies will not necessarily be those that get adopted at scale later and it is entirely unclear if, how and how much distributed ledger technologies will become integrated into economic sectors, including transport, and ultimately, into daily life. Uptake hinges on whether or not decentralised ledger technologies such as blockchain can deliver better value than traditional ledger and transaction frameworks in use today and whether or not they can enable new, value-adding applications that are not yet possible with existing technologies. This report explores where distributed ledger technology and associated changes in data syntax and inference may matter for people and discusses what anticipatory actions can be taken now to ensure that these developments support desired public policy outcomes.

The scope of transport activity is broad and intertwined covering both passenger and goods transport at local, regional, national and international scales. There are similarities between all of these when considering the impact that advances in data science will have but for the sake of this report, the focus will be on urban passenger transport. Where relevant, we draw in parallels from freight activity and intra-urban and international transport.

The transport sector is undeniably undergoing significant changes that have broad implications for people's everyday lives. These changes impact not only the set of "hardware" assets such as the vehicles and infrastructure that enable transport activity, but, also the data and "software" ecosystem that underpins transport activity. While the focus of transport and innovation policy has largely been on the former, this report investigates the latter. In particular, it describes the implications of recent advances in data science to the way in which transport services are delivered and, more broadly, the way in which transport activity is organised and regulated.

The interdisciplinary field of data science brings together scientific methods, processes, and systems to structure and extract insights and knowledge from various forms of data. This report focuses in particular on the convergence of data syntax, novel ways of structuring data – e.g. via distributed ledgers, such as blockchains - and methods for extracting insights from distributed data. All of these form a co-joined data ecosystem that supports transport activity.

Figure 1: **Schematic overview over data science**

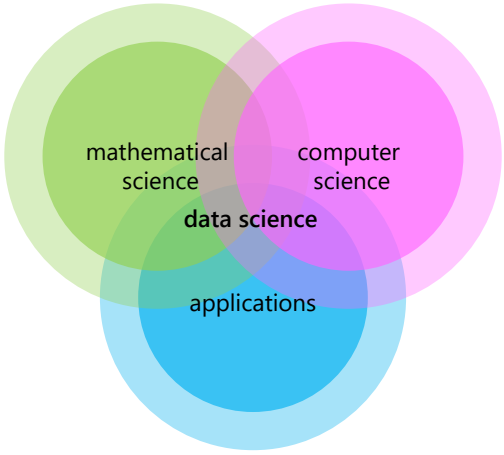

Source: adapted from (NYU, 2018)

While technology-led innovation will impact the deployment of new types of transportation assets and enable novel value propositions, the way the underlying data ecosystem develops will have deep and long-lasting implications on how transport contributes to individual and societal welfare.

The changes that the transport sector is undergoing are not isolated – the past few years have seen rapid and co-synchronous developments in technology, digitalisation, disintermediation, automation, changes in data production and advances in artificial intelligence and data science. All of these have significant implications taken in isolation. Taken as a whole and across sectors of human activity, they have the potential to change the way in which people and goods move in profound and difficult-to-predict ways. These developments are often spearheaded in the private sector and could improve several public policy outcomes, such as reducing congestion, increasing efficiency, reducing or removing environmental and social dis-benefits, or, they could do the opposite.

Uncertainty regarding the impacts of new technologies and services will be challenging for public authorities to address and manage but it seems clear that policy will have a role in guiding outcomes just as it seems clear that transport will increasingly be as much about bytes as about vehicles and infrastructure. Yet the scope for effective regulation and oversight has eroded as much of the information necessary to accomplish both tasks has shifted away from public authorities and regulators to the private sector. A further complication is that much of the data now held by the private sector is spread across multiple disparate and oftentimes competing entities.

New technologies and services have given consumers, including transport users greater agency to manage, source and curate their own travel experiences in the face of a broad range of commercial and public service providers and infrastructure managers. This meta-trend placing people at the core of business and public value propositions is one that also has long-term implications for transport governance and public policy.

Every trip starts with a simple desire: "Get me from point A to point B". Fulfilling that desire, however, may be quite complex if all available means of transport were to be considered and combined most efficiently. Doing that would require considering which means of transport are available, which assets have capacity, how fulfilling one trip might impact the fulfilment of further trips, how to plan, book, dispatch and access assets, how to coordinate transfers between assets and services and how to pay for services. The cumulative cognitive load of all these functions explains why consumers rarely take advantage of all mobility options available to them. But each cognitive "pain point" also reveals an opportunity for relieving that load through technical or organisational means.

Transport has largely been seen as a question of hardware – of vehicles and infrastructure. But transport is not about moving vehicles for the sake of moving vehicles, nor is it about building infrastructure for the sake of building infrastructure. Transport, and more precisely, the mobility it makes possible, delivers value by enabling people to access opportunities they cannot realise without movement, such as socialising, travel to work or education and access to services and goods. Realising these opportunities through transport relies on a web of interactions and transactions that run in the background of our daily lives but without which little would happen. Additionally, all of these trips rely on a large amount of data and information that build trust and enable these transactions to take place in a predictable and efficient manner.

Mobility requires infrastructure and vehicles but mismatches in supply and demand lead to over- and under-use of available resources and capacity. This happens despite the existence of information about transport options and, indeed, of information about activities that people are trying to access. Bridging the gap between instantaneous travel demand and transport supply will require access to data that is exceedingly personal (e.g. such as location, pattern of daily trips, inferred trip purposes, comfort/price preferences, etc.) and could potentially be subject to misuse. Nonetheless, individuals and society at large could reap considerable benefits if supplier-agnostic, user-centric, seamless mobility experiences could be delivered at scale.

This potential has important implications for policy since, over the last century, people have tended to opt for car-based transport whenever available, affordable and made practical. This is because car-use has generally responded well to peoples' desire for seamless, convenient and comfortable travel across a broad range of distances and in many urban contexts. Further, the affordability of car use has grown alongside growth in incomes and lower relative travel costs. But this growth has come at a cost to cities, people and society (e.g. congestion, unreliability of travel time, air pollution, crashes, car-dependency, inequitable access) that have eroded the benefits cars provide. This growing tension between cars and mobility in cities has led cities, citizens and companies to explore ways in which the benefits of car-like mobility can better be delivered across a wide range of mobility options by leveraging digital assets and data science. At the heart of these efforts are the many ways in which private and public actors are seeking to offer a seamless and rich Mobility as a Service (MaaS) offer that could compete effectively with – or integrate – private car-based mobility.

Mobility as a Service in a networked and meshed world

Urban mobility today is typically provided by a patchwork of poorly optimised and disconnected service providers operating with little coordination on both public and private infrastructure. Operators and public authorities have sought to optimise efficiencies for each mode but the combined effect of these efforts is still sub-optimal from an overall system efficiency perspective.

At any given time, even at peak periods, cities are flush with unused transport capacity. There are many reasons for this. Information about available services is poorly distributed across the travelling public. People may also choose to ignore modes with available capacity if they feel these to be unreliable or undesirable. Finally much available capacity remains unused since it is often scaled for peak demand and no market yet exists for off-peak uses of much of it.

Today, however, choices available to citizens are expanding and changing rapidly in an ever more dynamic urban mobility ecosystem. Traditional stakeholders including public transport operators, car manufacturers and taxi companies are facing increasing pressure to innovate, to attract and retain users and, to do so, have to alternatively compete and co-operate with each other and new market entrants proposing novel business models.

Against this backdrop is the consumer who just wants to get from point A to point B in the most demand-responsive, flexible, pain-free, reliable and affordable way. As in other areas of their lives, they want trip experiences that place them in control and which leverage the most convenient options available irrespective of who offers them. And the number of actors offering new transport services is growing as technology opens up new possibilities in accessing shared resources, automating vehicle systems and connecting supply and demand. Whereas transport has been a siloed world of independent and separately regulated services (Figure 2), the future of urban mobility may very well be much more aligned with other "as-a-service" models where actors engage directly with each other on the basis of commonly agreed protocols in a much more mesh-y world.

What is Mobility as a Service?

Mobility as a Service (MaaS) or Transportation as a Service in North America) is a term used interchangeably to describe packages of bundled transport services or, more generally, as a broad concept describing new, customer-centric ways of seamlessly accessing a range of different transport services – many of them shared. There is no single definition of MaaS but the UK Transport Systems Catapult offers one that highlights its key features:

"[MaaS uses] a digital interface to source and manage the provision of a transport-related service(s) which meets the mobility requirements of a customer" (Transport Systems Catapult, 2016)

MaaS represents a break with the past in that mobility services have historically been provided by siloed operators, manufacturers and public authorities with little practical cross-mode coordination. While some public transport operators have sought to provide a more diversified offer and some other actors (including most recently, vehicle manufacturers) have sought to offer more flexible ways to access cars, scooters and bicycles, there has been little real joining up of legacy and emerging services into a simple, open, customer-interface.

At its core, the concept of MaaS supports the digital joining-up of different transport, information and payment services into a smooth and reliable customer-facing experience (Figure 3). These services may be those provided by a single operator in cases where extensive integration exists or may involve a MaaS provider bringing together services offered by third parties into a coherent framework. MaaS supports the integration of public transport modes, commercial transport services such as ride services, bike and ride

sharing and taxis into a comprehensive mobility offer. By providing a smooth, convenient and dependable travel option across multiple operators, MaaS could favourably compete with individual car use for some people and enable more efficient use of transport assets. Indeed, proponents of MaaS would see vehicle ownership becoming secondary to, and less attractive than, MaaS.

Figure 2: **Siloed mobility services**

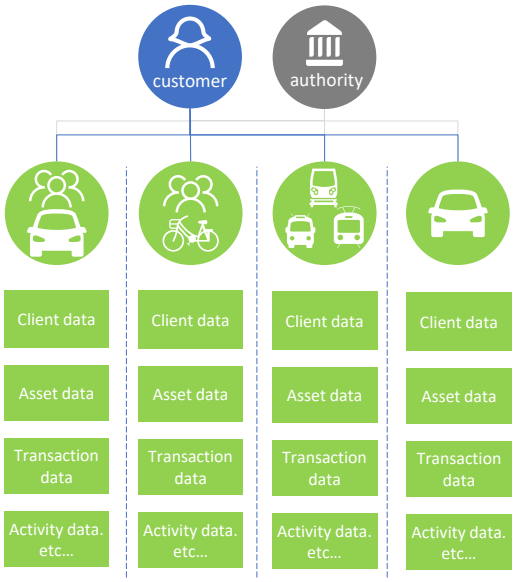

Figure 3**: Mobility as a Service leverages digitalisation for customer-centric transport services**

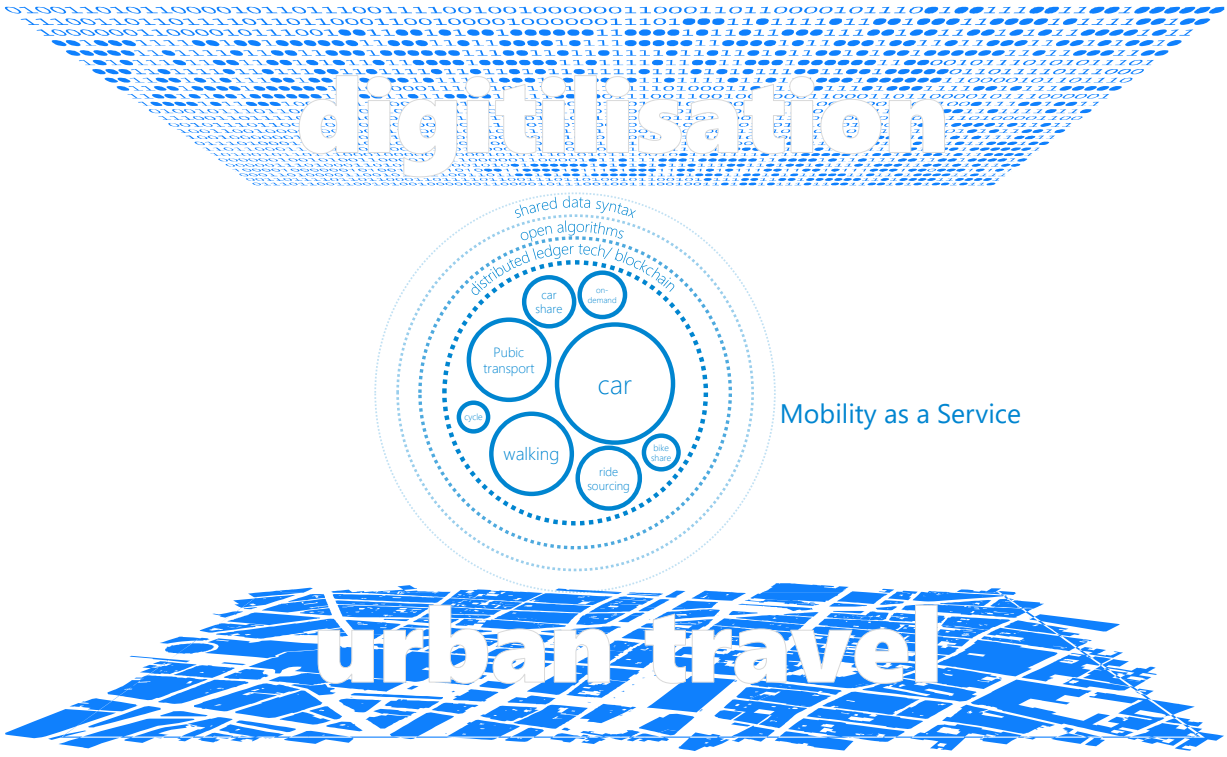

MaaS can be tailored to individuals' needs, budgets and constraints. Its delivery involves identifying clients and operators, gathering information about availability of services and capacity, and managing payment and revenue allocation within a common framework. In some models, customer-facing MaaS may take the form of a subscription to a pre-negotiated package or bundle of services like those offered by the Whim App developed by MaaS Global (https://maas.global). The offer may also be structured along a "pay-as-you-go" model that coordinates services and payment within a common customer-facing environment. Both approaches may also co-exist in the same model. Multiple hybrid forms of commercial MaaS-like services have been announced by companies such Ford, Transdev, Didi Chuxing, Uber, Moovel, Fluidtime, Waymo, etc. Some of these are more limited in their scope whereas others –through their coverage, number of mobility services offered and stated goals, are more ambitious.

In practice, there is a continuum of MaaS-like arrangements that extend from single-operators offering multiple services, to an all-encompassing MaaS platform that federates different and independent transport service providers. On the single-operator side of the spectrum are entities that provide vertically integrated services. These might include public transport operators that provide both bus and rail-based services (and shared bicycles as some do) or a commercial operator that provides different classes of taxi or app-based ride services. The other end of the spectrum is currently unexplored territory as there are no cases of a single operational platform that federates all transport service providers within an integrated and seamless framework. What can be seen is a generalised move from the former towards the latter along a number of different trajectories.

Multiple iterations and implementations of MaaS-like arrangements exist. All of these are vulnerable to being disrupted by the widespread deployment of decentralised applications to varying extents. To understand how and where, it may first be helpful to decompose MaaS into its principal components, typical stakeholders and business processes.

Building blocks for Mobility as a Service

In its broadest and potentially most compelling implementation, MaaS requires several components. It first requires physical transport and infrastructure assets. The costs for deploying, maintaining, renewing and building these must be borne by commercial operators or by public authorities in an environment where the roles of both in providing and overseeing transport activity are shifting. A key challenge in the deployment of more integrated forms of MaaS is the financial model that allows those investing in infrastructure to cover the costs of doing so. This has been a traditional barrier for public transport operators to opening their data and allowing it to be integrated with third-party platforms. Another roadblock to further integration of services, especially those operated or controlled by public authorities, are administrative barriers inherent to separately managed departments – including those run under different administrative regimes.

Additionally, public transport operators often invest considerable resources in branding and marketing their identity as a way of retaining and increasing ridership and it isn't clear if integrating their services in a broader MaaS offer would deter or boost these efforts. In any case, this concern highlights that interests may not be uniformly aligned in a broad MaaS ecosystem and this has likely been a barrier in wider uptake (Polis, 2017).

Beyond transport, MaaS requires an installed base of assets and infrastructure that ensure digital connectivity (3G-xG, WIFI, RFID, microwave, etc.) as well as the set of hardware devices and operating systems that allow customers to access services. These assets and infrastructure are managed outside of the transport sector by information and communications technology (ICT) industry actors and protocols.

MaaS requires a set of transparent, vetted and trusted commercial agreements that encompass commercial operators, public services and third-party aggregators of services (where applicable) and should cover

payment and revenue allocation amongst all parties. These agreements should enable viable services to be developed by all parties. At the same time, these arrangements should meet market power tests to ensure that undue concentration does not lead to an erosion of consumer welfare and inequitable service delivery. These commercial agreements will cover services that may be complementary to each other but in some cases, may lead to competition amongst different parties for some types of activities. Negotiating these will likely require all parties to adjust expectations and traditional service delivery models.

MaaS also requires open information about transport services. This information can be mediated and linked by third-party applications and way-finding services like those integrated into several search portals or operating platforms (Google-Android, Apple-IOS, Baidu, etc.). Alternatively, this information can be curated in third-party wayfinding applications like CityMapper, Moovel, Navtime, or Moovit. Information about transport services can be merged on closed platforms (as single-integrated models provide – this is the case with some ride service operators and many public transport operators) or an open platform where data on different services and offers are integrated. Various hybrid implementations of these information access methods also exist and are likely to develop over time.

MaaS will also have to find a market – this market will depend on varying local contexts, and the expectations and experiences of potential clients. The market will also depend on a set of policies that may facilitate the uptake of MaaS (e.g. pricing parking or road use, open data requirements, zoning that favours high density use, fiscal disincentivisation of company cars, etc.) or work against it (preventing sharing of data by public transport operators, banning new mobility service providers, under-funding maintenance of public transport infrastructure, etc.).

Stakeholders in Mobility as a Service

The delivery of comprehensive MaaS involves several stakeholder types (Transport Systems Catapult, 2016).At the centre of the MaaS ecosystem is the customer who has a desire to go from point A to point B, and who would like to do this in a seamless, convenient and affordable way (Figure 4). Travel decisions are not developing in isolation of other decisions or trends in society. Many people are changing the way they organise other activities (e.g. shopping, telephone and internet subscriptions, use of software, access to cultural goods such as music and other media, insurance, etc.) thanks to the technology-enabled and platform-facilitated disintermediation of consumption.

This broader trend towards "service-itisation" has allowed consumers to bypass traditional middlemen and connect to markets directly (e.g. Amazon.com and AliExpress for shopping) or create entirely new markets (e.g. Airbnb for short-term lodging or BlaBlaCar for spare carpooling capacity). People have growing expectations that they should be able to experience single-point-of-interface and seamless experiences for their mobility choices as well.

Transport operators and vehicle resellers (for individual motorised transport) have been the traditional point of interface for people regarding travel options. They have also been the principal point of regulatory focus with each category of activity (car use, public transport, regional rail, informal transport, taxis, etc.) operating under a separate regulatory framework. This set of stakeholders has been relatively stable over the past century representing private motorised and non-motorised modes and public transport comprised of bus, rail and subway services. Novel technology-driven transport services including car-, bicycle- and scooter-sharing, ride services, new forms of on-demand micro-transit and crowd-sourced transport services have disrupted these relatively stable markets.

Figure 4: **The Mobility as a Service ecosystem**

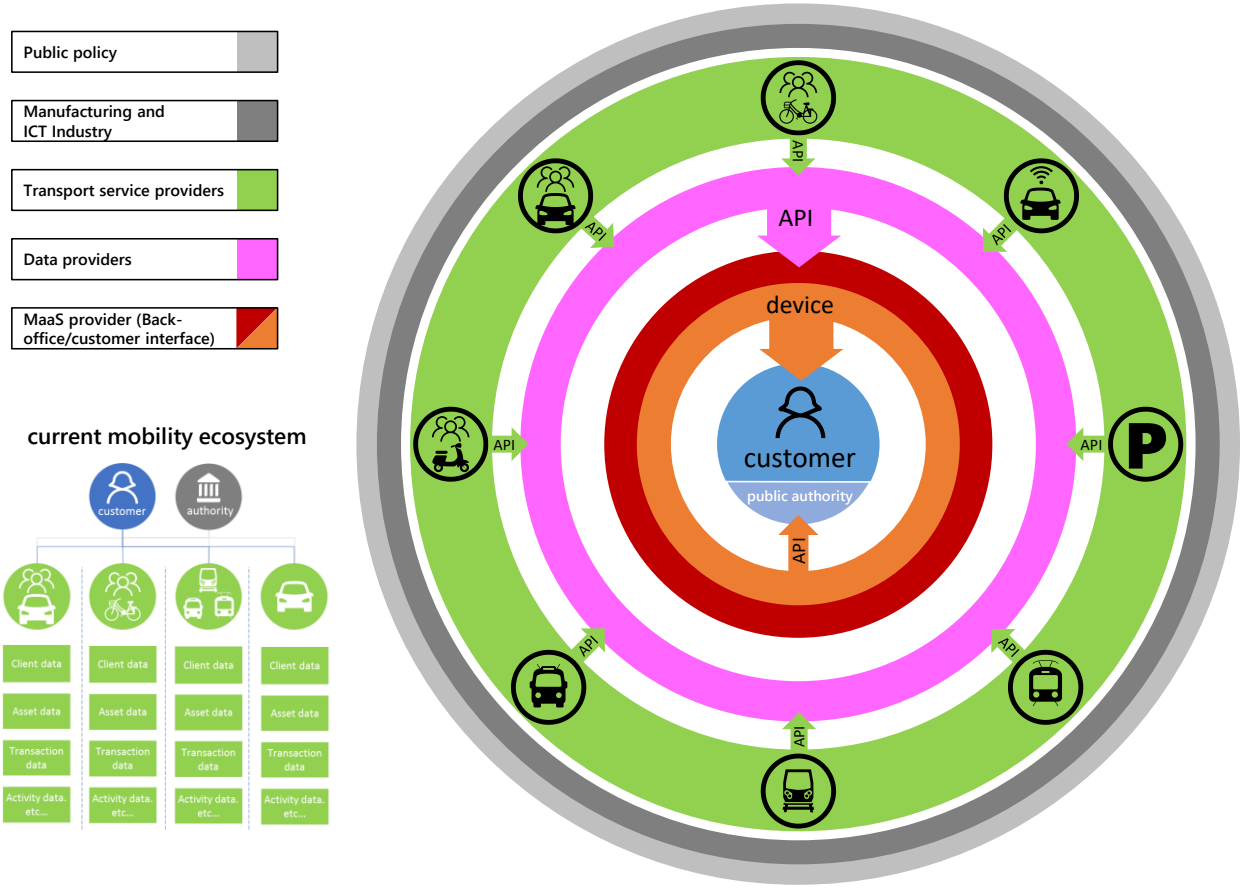

Data providers compile information on transport services, schedules, destinations, incidents and, to a certain extent, can directly facilitate links between different operators. They do this by accessing information manually from transport operators or by accessing dedicated open- or permissioned data portals put in place by transport operators. The latter provide Application Programming Interface (API) access to a sub-set of their data (either scheduled or real-time) so that it may be included in third-party applications. Operators may also provide software development kits (SDKs) so that developers can build their own compatible software that exploits transport operator data. In the case where operators do not provide data in digital form, third party data-providers may manually transcribe published schedule and service data into machine-readable form though this may be of questionable legality in some jurisdictions.

API-sourced or other machine-readable data have various formats and compiling them across a wide range of providers into a single operational environment for MaaS requires automated data-translation routines which must be adjusted for changes in operator data structure. This work of data translation and cross-platform harmonization is at the core of many mapping and navigation service business models. Data providers have in many cases disrupted legacy business models that exploited large information asymmetries between consumers and transport operators, but they too are vulnerable to disruption, potentially from the widespread uptake of a common data syntax for mobility that would enable data to speak to data directly without having to go through third-party "translation".

In a fully built-out MaaS ecosystem, the MaaS provider is the entity that makes the link between various transport operators and individual customers building on information provided by data providers. It is the virtual "agent" that collects information on where, when and how travellers want to move, pools information

on available transport capacity and price, negotiates packages or fares, provides routing and trip-making information handles post-trip clearing operations. The principal innovation here is the skill with which the Maas provider aggregates transport operator services on a digital platform, compiles and processes them in a unified environment and provides consumers with a value proposition for seamless trips that satisfies their demands. MaaS putatively broadens the mobility market by making travel more convenient and thus may create new business opportunities and yet engaging transport operators to participate in a common MaaS platform where they may compete with each other for some trips is challenging in many instances. While consumers may want seamless access to a wide range of services, it is not clear that all operators are ready or willing to engage with consumers in a completely open MaaS environment.

There are two other stakeholders of relevance to the MaaS vision. The first is the set of manufacturing and technology companies that provide the material basis for MaaS. These include manufacturers that provide cars, mini-buses, bicycles, scooters, sensors, radio equipment and other hardware that are essential for transport operators' services. Many of these manufacturers are introducing higher levels of automated driving for their fleets and in some instances, are designing vehicles that can be fully integrated into MaaS ecosystems. These manufacturers also include companies building smartphone and other interface devices that consumers use to access transport services. Secondly, the communications operators that provide the connectivity required for ubiquitous MaaS operations are part of the broad set of technology facilitators for MaaS.

Final principal actors in the MaaS ecosystem are the public authorities who set the regulatory framework for transport and communications. These authorities have typically addressed each individual mode of transport separately and have been challenged by the increasing hybridisation of these in ways that were not anticipated. Some countries, such as Finland (Box 3) have sought to reset their transport legislation to support the broad implementation of MaaS-like services but this remains more the exception than the rule. Broader uptake of MaaS will require revisiting regulations and, in some cases, putting in place a set of rules for MaaS providers.

Essential functions and processes in Mobility as a Service:

MaaS ecosystems rely on a number of processes that ensure a seamless trip-making experience from the customer's perspective. These processes occur irrespective of whether MaaS is delivered by one or several operators or providers. For each of these processes there are a series of corresponding technical methods that support them. While these methods are undergoing considerable flux as database, identity management, data access and transmission protocols co-evolve, they are largely based on permissioned- and API-mediated access to in-house or cloud-based proprietary databases.

Secure identity and access management: The identity of users, operators, service providers must be established in a trustworthy manner and this identity must be linked to rights to use services (and thus linked to payment data) or to dispense services (and thus linked to certification and licensing).

Authentication: The identity of users and service providers must be authenticated across multiple services and multiple use cases.

Asset identification: Assets should be identified and data related to them authenticated. Available capacity, location, vehicle condition and type, state of repair, etc. should be discoverable to all processes seeking to fulfil relevant user trip requests.

Service specification: Fulfilling MaaS requests requires cross-platform and easily accessible information about available service types. These may include on-demand operation, station or stationless sharing, scheduled services, shared versus exclusive use, different service classes, etc.

Routing and connection information: At the heart of MaaS are the back-office mechanisms that join-up different services within or across operators so that travellers experience seamless trips. These mechanisms combine real-time routing and, if necessary, connection information so that people can reliably switch from service to service or from mode to mode as if they were just one.

Near real-time access to information: Asset, routing and connection information should be accessible in as close to real-time as possible so as to reflect the actual trip-making environment accounting for changes in traffic, off-schedule operation or other factors that might impact the reliability of travel.

Transaction processing and clearing mechanisms: Users accessing services across multiple providers require some form of commonly agreed booking, invoicing, processing and clearing mechanism to ensure that rights are matched to users as they switch from one operator to the next. At the same time, revenue allocation mechanisms must address how operators are to be compensated for their fractional contribution to a total trip chain. These mechanisms should allow all parties to achieve consensus on what resources were used to fulfil a trip and how payments for these were allocated across all actors.

Payment mechanisms: The actual payment mechanism should allow for seamless and unitary payment for services from the customer's perspective and should be tied into the back-office transaction and clearing mechanisms.

Data logging/sharing and transmission: Data generated by sensor platforms and embarked on vehicles or devices carried by people, and transaction and trip-related data all underpin the delivery of MaaS services. This data is necessary for delivering real-time and high-quality user experiences. In aggregate form, it can also help deliver better overall transport system performance. MaaS operators and providers record this information and either make some or, more rarely, all of it available for use by others in the ecosystem. Data access rules are typically set up on a case-by-case basis as much of the data is commercially valuable and could prove to be invasive to privacy.

Efficient and secure distribution of information: Data on transactions and trips is the lifeblood of the MaaS ecosystem. A data sharing framework that quickly and efficiently allows the cross-platform sharing of relevant and timely information is a core requirement for MaaS.

Disintermediation: MaaS seeks to digitally streamline the joining-up of different transport service providers with customers. Although it may introduce a new intermediary in the customer-service provider relationship – the MaaS provider or platform – it at the same time seeks to simplify that relationship from one-to-many to a one-to-one relationship from the user perspective. That disintermediation is analogous to similar trends occurring in other sectors.

Mobility as a Service in a networked and meshed world

Central to the "as-a-service" concept is the platform model that helps to federate stakeholders and services. Though platforms can be fully open to participants, many existing implementations of MaaS conform to a hierarchical and closed model in which interactions between different actors are fixed in pre-determined, negotiated relationships and are open only to vetted and centrally-permissioned participants.

The potential positioning of MaaS along two axes – one running from centralised to decentralised models and the other running from closed to open systems – helps to clarify where advances in data science create new possibilities for MaaS (Figure 5) (Casey & Valovirta, 2016).

Figure 5: **States of the Market Value Systems of Mobility as a Service**

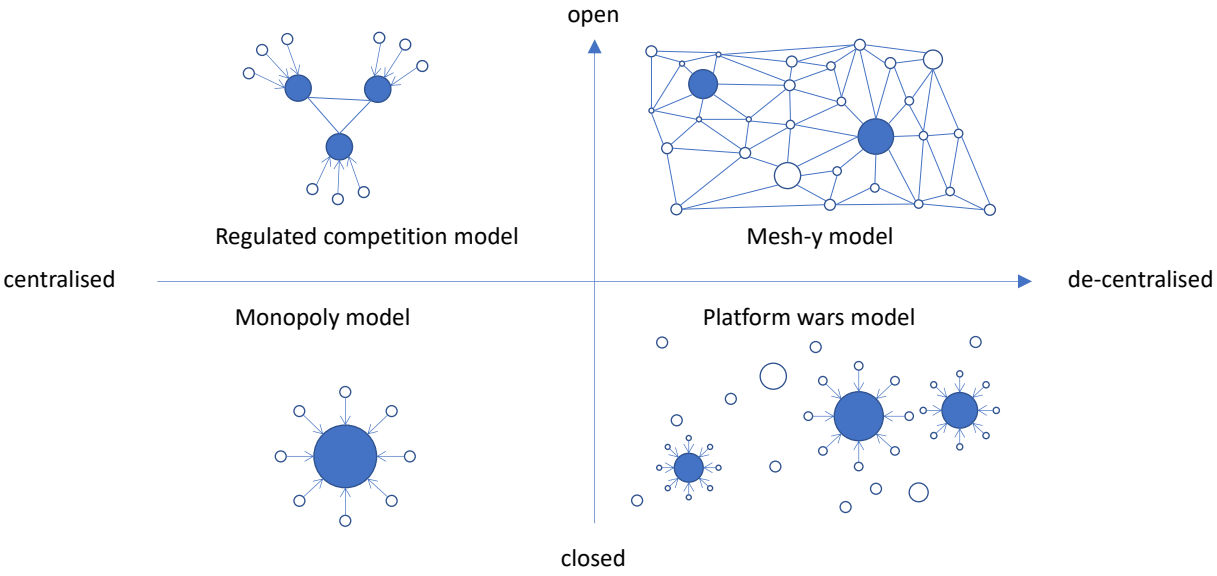

Source: Adapted from (Casey & Valovirta, 2016)

Today, many transport services and their regulations tend to be aligned with centralised (and often public) actors. Provision and maintenance of road and rail infrastructure, planning, driver licensing, safety regulation, vehicle registration and taxation, public transport services, taxis and para-transit all involve centrally-provided services or government regulation. The centralised nature of these services is a result of their displaying strong public goods characteristics. Public goods are those that, because they are indivisible and non-exclusionary, may not be satisfactorily or equitably provided by market actors. Centralised and regulated models may deliver guaranteed service quality, but this may come at the price of lower innovation. New technologies, business models and novel regulatory structures may open up possibilities for these previously indivisible and non-exclusionary services to be satisfactorily delivered in a decentralised context while retaining the quality control aspects of centralised systems.

At the same time, a significant part of the transport system can also be seen as being decentralised, that is under the control of, and operated by private actors. For example, a significant share of vehicle parking capacity is private and most traffic on roads is the outcome of households and commercial actors making largely uncoordinated (but not unpredictable) travel decisions. Transport operators may deploy IT systems to optimise their own operations and asset deployment just as households may seek to optimise their own trips using third-party way-finding and navigation applications, but these actions are rarely explicitly coordinated with others.

Closed systems can ensure predictable service delivery and low transaction costs since all services are coordinated within a constrained set of actors and/or functions. The model of the centralised and closed "in-house" MaaS platform (or MaaS-like platform) is a compelling one in that the MaaS provider can offer a coordinated and highly customised user experience (Summermann, Oge, Smolenski, Fridgen, & Rieger, 2017).

Centralised and in-house platforms collect data and gain knowledge about users that enable companies to provide innovative and valuable services. As an example, the emergence of successful ride-service platforms like those of Didi Chuxing, Uber, Lyft, Grab and Ola has been fuelled not only by the attraction of their core ride service product but also because of constant service innovation (e.g. shared services, on-demand minibuses, links to bicycle sharing or public transport) whose combined offer begins to

approximate a full implementation of MaaS. Other multi-actor MaaS platform models, like those of MaaS Global and Ford seek to offer a more open ecosystem where the MaaS provider coordinates, but does not develop nor control, the core transport services made available to customers.

Open systems, on the other hand, are less likely to create lock-in effects and information monopolies, allow for a higher level of data control and ownership for individuals, facilitate sharing of resources and assets and allow for the co-creation of innovative business solutions and partnerships (Summermann, Oge, Smolenski, Fridgen, & Rieger, 2017).

Platforms derive their principal strength from the network effects of connecting people and markets through the use of digital infrastructure. Platforms are sensitive to network externalities in that the value of goods and services traded increases with the number participants involved. Few people, for example, would be interested in joining a MaaS platform if it only offered extremely limited services just as few transport operators would wish to join a platform if it had a small customer base with no growth prospects.

Platforms face a challenge in that they must simultaneously and rapidly attract customers and service providers to reap the benefits of scaled-up network effects. This has been a defining feature of the early development of ride-service platforms in that they have had to attract and retain drivers and other service providers in order to, in turn, attract and retain clients which then make the platforms attractive to more drivers and service providers. Though the multiple platform model currently is the emerging paradigm for MaaS, it is not the only market configuration for consumers and cities.

Four possible mobility value system states emerge when looking at the intersection of centrality and openness (Figure 5).

In the lower left-hand quadrant is the closed and centralised world of the single monopoly platform or service. It is a world in which one actor has out-competed all others or has been installed as the single MaaS or transport operator. According to this model, the monopoly actor organises, controls and delivers service(s) via vertically integrated and closed technical systems and asset bases (vehicles, payment, routing and dispatching, information systems, etc.). This quadrant is fraught with the risk of anti-competitive behaviour which requires a commensurate and rigid regulatory response from authorities. Monopolistic actors are typically reticent to new market entry, have a poor innovation track record and are slow to recognise and adapt to external changes. While monopolistic market domination is a naturally attractive model for businesses, it is one that is unlikely to deliver robust or durable societal benefits especially in the face of rapid societal and technological changes.

In the lower right-hand quadrant is the world in which multiple isolated platforms or services compete with each other with little or no coordination or cooperation amongst themselves. Market actors deploy proprietary and incompatible technology and business systems leading to a fragmented mobility service market. This "battle of the platforms" world is the one that is best characterised by the current state of MaaS-like implementations. It also describes the current state of play for many cities where transport operators operate and are regulated in rigid silos.

Platforms and transport operators competing against each other is not necessarily unhealthy per se in that innovation and lower prices may result, but anti-competitive behaviour may also emerge if homing costs (costs associated with affiliating to a platform) and switching costs (costs associated with moving to another platform) are high and thus contribute to lock-in effects. These would result in a shift from this model towards the monopoly model. While direct homing and switching costs are generally low for MaaS platform participants, there are many indirect costs which may contribute to lock-in effects. These may relate to contractual terms, pre-paid and non-portable subscriptions, the non-portability of personal data, platform-specific investment in assets and the potential opportunity cost of switching from one platform to another when platform growth trajectories are uncertain.

Multiple platforms competing against each other may also lead to sub-optimal uses of overall resources, especially in the absence of externality pricing, since each platform or operator will deploy or mobilise duplicative assets and services to compete for consumers. The overall balance in a multi-platform world between upside service innovation and lower prices and the downside potential for lock-in and reduced overall system efficiency is not clear and deserves attention since more and more transport services will likely be delivered via MaaS-like platforms.

In the upper left-hand quadrant is a partially open and centrally-regulated market with a few, coordinated yet competing actors and a common set of technical and operating standards. This regulated competition model is what can be seen in cellular communications markets where competition for market entry does not exclude cooperation amongst actors for certain service components (e.g. technical standards for equipment interoperability, data portability requirements, roaming charges, etc.). Consumers can easily switch among service providers and platforms in this model, but the number of market actors is limited by public authorities.

Finally, the upper right-hand quadrant describes a world with multiple, loosely-coordinated market actors, operating in a space built on standardised data and technical interfaces as well as common data exchange and processing protocols. Switching and homing costs are low or inexistent and the openness of the system leads to rapid and customer-centric innovation. This is a much more "mesh-y" model than the others and one where the technical means of service production, distribution and transaction-clearing are distributed and democratised. It is roughly analogous to the open internet model that is characterised by a great diversity of actors and business models operating on top of a shared Transmission Control Protocol/Internet Protocol (TCP/IP). Without centralised control, it may seem that the quality of services delivered in this quadrant may be more difficult to control but this is precisely where recent advances in data science may help.

In an everything-to-everything meshed world, a diverse set of stakeholders must trust each other - from start-ups to established companies, individuals to governments, both within and outside the transport sector. In order to willingly and freely initiate transactions amongst themselves, they must also trust underlying business processes and those relating to service delivery, payment and clearing operations.

Just as the internet established trust in the ability to quickly, accurately and predictably communicate information over a web of connected machines, new developments in data protocols, structure and data science may help deliver value by robustly establishing this trust.

"Everything to everything" interoperability

Delivering on the promise of MaaS in a meshed world of customers, transport operators, data providers, infrastructure and asset owners, and public authorities will require a significant shift away from the *status quo* covering existing protocols and business logic. Ensuring this shift enables "everything-to-everything" interoperability in the context of MaaS will involve action in the following three areas:

- **Distributed Ledger Technology** such as blockchain: A move away from platform-based MaaS frameworks to those where markets for supply and demand are cleared near-instantaneously and with little centralised control may open up the potential for widespread uptake of MaaS. Central to these frameworks is the way in which distributed and non-centralised trust, robust identification and authentication functions and transaction clearing are carried out.
- **Data syntax** for MaaS: A common data syntax for encoding the various components of MaaS would facilitate the uptake of these services. Currently, the lack of common data structures between different commercial and public transport services serves as a barrier to their integration.
- **Open Algorithms** and other alternatives to data-sharing: MaaS requires commercial partners to share a considerable amount of data either amongst themselves or with centralised platforms. This sharing of data is often problematic for operators who view their data as commercially sensitive and privacy-relevant for their clients. Alternatives to data sharing that enable stakeholders to access vetted, trusted and actionable insight from proprietary data may remove many roadblocks to broad MaaS partnerships.

Blockchain and distributed ledger technologies

Blockchain and other forms of Distributed Ledger Technologies (DLTs – e.g. Ethereum, or other privately developed systems that use parts of the blockchain concept to establish a network of trusted nodes) combine recent advances in data science, cryptography, and novel governance principles – and have been highlighted as one of the most disruptive sets of technologies since the advent of the internet. Blockchain is

> *"… neither an innovation per se, nor an object, but rather the intelligent and unprecedented combination, variable according to different actors, services and existing technology platforms, to create a "collaborative management of a distributed registry"– a system to create trust between actors without resorting to centralised governance and organization invested with wide-ranging and exorbitant powers. In the world of trust and of databases, both marked by decades of stable models and thinking, this change is radical. It falls outside of current skillsets, legal and regulatory frameworks, business models and installed IT systems, software and infrastructure." (Dardayrol, 2017)*

DLTs such as blockchain are general-purpose tools that are characterised by four fundamental attributes: transparency, traceability, trust and distributed governance. They have the potential to create new business and regulatory processes by:

- authenticating ownership and rights and ensuring secure value transfer across a wide range of stakeholders and assets;
- creating an immutable, distributed and openly verifiable record of past transactions;
- being designed specifically for secure data transfer and parametric data privacy/access with built-in encryption protocols;

- ensuring robust cyber-security via the redundant and secure nature of transaction-validating nodes.

DLTs also foster innovation and efficiency via automated business or regulatory processes via condition-dependent, self-executing algorithms known as "smart contracts".

Although different from previous innovations that have led to the development of the internet and its associated value chain, DLTs build on these. The internet's strength has been the development of a scalable and reliable shared protocol for exchanging data amongst a network of connected computers and other devices. DLTs enable the trusted exchange of value and transactional information leveraging distributed and authenticated ledgers (Figure 6). In this way, DLTs can be the basis for implementing "mesh-y" MaaS (Figure 7).

Figure 6: **Internet and distributed ledger technologies compared**

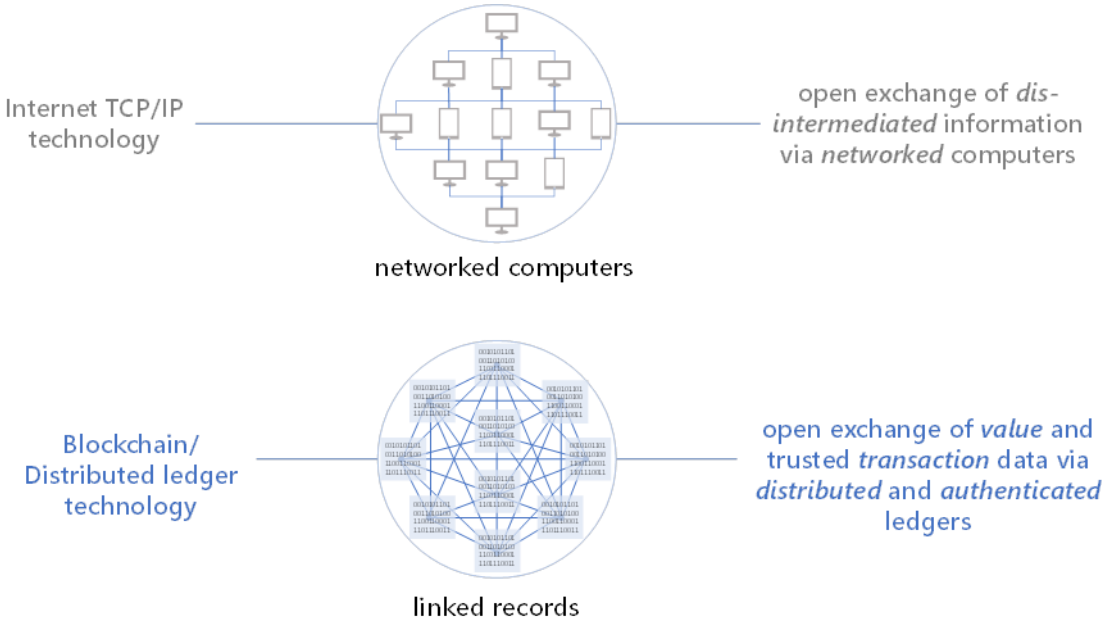

Figure 7: **How distributed ledgers and blockchain enable Mobility as a Service in a "mesh-y" world**

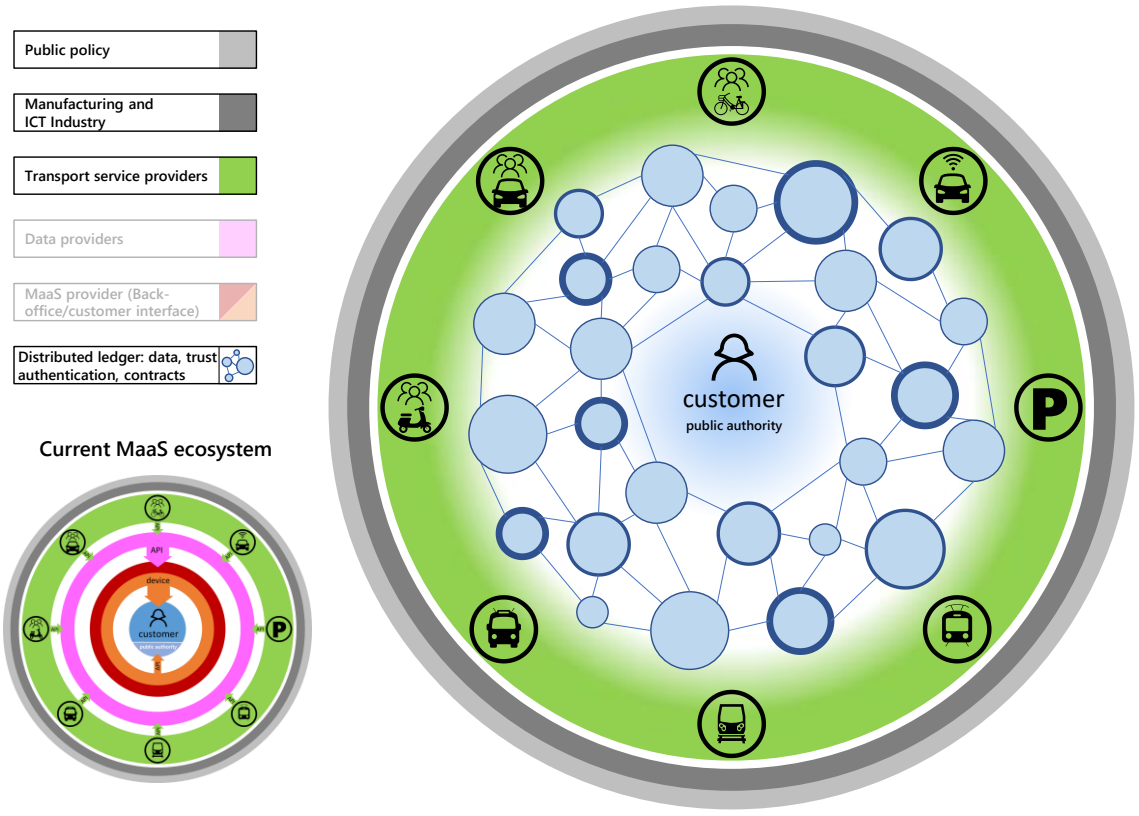

At the time of its conception in 2008, blockchain technology was intended as the underlying method to create and safeguard the value of a new currency: Bitcoins. Since then its core concepts have been transferred to other cryptocurrencies and onward to wider applications. What all of these applications share in common is that they must address how to enable frictionless transactions regarding ID, authentication and access to distributed services, data and rights without the need for a third party to establish trust between transacting parties. Blockchain and its early implementations hint at a broad set of uses – including within transport and in support of MaaS applications – but scalable examples are still in early stages and few robust applications have been developed outside of the digital currency domain.

The contribution of blockchains and DLTs to enabling a more open and seamless MaaS ecosystem will depend on how well the technology can scale to match potential use cases, how well it performs compared to existing processes and arrangements and, ultimately, how broadly it is adopted by consumers, businesses and public authorities. The outcomes of these factors are linked to the technology itself and its fitness for purpose for MaaS.

Blockchain and transport

Blockchain and other distributed ledger technologies create new possibilities to manage distributed and fractional capacity (both for vehicles and infrastructure) and offer the possibility for customised, dynamic and sized-for-purpose transport to individuals. It also potentially allows operators to manage access rights, data and payments across a broad network of unrelated and competing transport service providers and platforms. Indeed, much in the MaaS ecosystem is "blockchain-able" (Table 1).

Table 1: **MaaS processes and distributed ledger technologies**

MaaS Process	Blockchain suitability
Secure identity and access management:	High: Private key identity management built into Blockchain protocols can support this function.
Authentication:	High: Consensus-backed authentication using private public keys supports this function.
Asset identification:	High: Immutability of records combined with strong ID and authentication protocols support this function.
Service specification:	Medium: The specification of services can be included in transaction block but having a common syntax for the Internet of mobility would be helpful here.
Routing and connection information:	Low to medium: Dynamic routing and connection information would require low latency that is challenging for many current blockchain applications. These functions could be run "off-chain" and securely referenced in the protocol.
Near real-time access to information:	Low: Same as above.
Transaction processing and clearing mechanisms:	High for high-latency operations (like subscribing a service), Low for real-time micropayments (Though new forms of DLTs may address this).
Payment mechanisms:	Medium: There is considerable uncertainty regarding various cryptocurrency models. Building these into Maas at this stage may be premature and risky.
Data logging/sharing and transmission	Low: current implementations of blockchain technology are poorly suited for real-time or high frequency logging. This may change with new generations of DLTs.
Efficient and secure distribution of information:	High: Distributed and secure data management is one of the foundational principles of blockchain technology.
Disintermediation:	Medium: In a fully built-out DLT supported MaaS ecosystem, many intermediaries may no longer be necessary. How well DLTs like Blockchain can automatically carry out the functions currently ensured by MaaS intermediaries like data suppliers, however, is still unclear.

Applications for passenger transport

Blockchain and other DLT technology could significantly open the scope of MaaS interoperability and optimisation by facilitating direct and platformless MaaS applications. DLT technology could improve passenger and service provider ID authentication within and across different transport modes to support the implementation of a MaaS across multiple public and private transport providers. DLT could also help facilitate multi-party payment and revenue allocation clearing functions and, when combined with "Smart" contracts, could enable direct, on-the-fly service and payment integration in a much more open framework than that provided by platforms.

Following the trend away from data processing in siloed databases and cloud applications and towards computation occurring directly at the point of value creation (analogous to edge computing for sensor-based data), DLT could deliver significant gains towards creating an "internet" of mobility. The exact potential for DLTs to contribute to MaaS and other applications in transport will be linked to the specific ways in which DLTs function differently from existing methods. In particular, this potential will be highly dependent on whether DLTs can provide better ways of doing things than existing approaches. This is not a given across the wide range of MaaS tasks.

At present, there are many early stage and exploratory initiatives by both start-ups and more established companies to leverage blockchain technology for transport-related applications. These initiatives are first steps to understand the technology, build capacity within the transport sector to understand the potential and the limits of DLTs, establish alliances to drive the development of blockchain applications according to specific transport sector needs and for start-ups to position themselves ahead of the competition by demonstrating solutions. It's important to note that much of what is being proposed in this field is either in initial whitepapers or limited proof-of-concept trials – there is little clarity on which applications, models or protocols will move beyond this stage.

The field is currently very volatile, with new players emerging on a weekly basis, while first movers like Arcade City and La'Zooz (both blockchain-powered peer-to-peer ride-sharing services) have faced difficulties and have since pivoted to other services or business models. Arcade city is now Swarm City, another peer-to-peer ride service concept powered by Ethereum (https://swarm.city/) and La'Zooz is no longer active with the original team now developing Commuterz (https://www.commuterz.io) a blockchain-powered carpooling service.

Toyota Research Institute has teamed up with MIT Media Lab and partners BigChainDB, Oaken Innovations and Commuterz to build a blockchain-powered car-based mobility ecosystem. At the outset, the partnership will focus on three areas. First, it will create a blockchain protocol for anonymously sharing safety performance-related data from automated and eventually, fully autonomous, cars – possibly with monetised incentives. Today, safety-related information (outside of crash-related data) is siloed by technology companies or car manufacturers creating a situation where the safety performance of vehicles may be uneven across the fleet. Pooling this data across multiple drivers, vehicle fleets and manufacturers should accelerate the learning curve regarding safety-critical performance across all vehicles – this is also an area where other DLT applications are being explored as in the case of a newly announced initiative by the IOTA foundation (described further on).

Second, the Toyota Research Foundation-MIT initiative will develop blockchain apps to support peer-to-peer transactions that allow vehicle owners to sell rides, cargo space or even rent the vehicle itself. Finally, the partnership will explore ways to develop blockchain applications in support of usage-based insurance that would leverage data from vehicle sensors and thus reward safe drivers with lower insurance fees.

These multi-pronged approaches are being practised by others as well. DOVU (DOVU, 2018), a Jaguar Land Rover-backed UK start-up, has a multi-layered approach that also includes the digitalisation of traditional business cases like the insurance of cars and drivers: for example, they explore how a blockchain-based data marketplace could combine car data with insurance data and driver history data to calculate smart pay-as-you-use insurance policies for the use of shared – or shared, automated cars in the future. Their model includes encoding this information in Ethereum smart contracts that will automatically execute claims once an insured event happens. They also foresee blockchain-based rewards for allowing people, firms or any "data owner" (which could potentially even be an automated vehicle providing rides) to monetise driving data that can then help with traffic planning and other uses.

DOVU's vision goes further still, including elements that could be considered part of a MaaS package. The DOVU ecosystem accounts for linkages and with a broad range of mobility providers, including public transport operators. It also describes how it's data marketplace could support sustainable transport and active mobility: e.g. walking or cycling (as tracked by smart phones and authenticated via a blockchain) will earn users tokens that then can be exchanged for reduced ticket costs for public transport or privileged access to a vehicle when needed.

Similar visions are shared by Ernst & Young, which launched their "Tesseract" platform (EY, 2017) - also in the UK) within weeks of DOVU. It uses blockchain to manage access to single vehicles, fleets and other transport services on a single platform by digitally logging all the information on the blockchain and automatically settling all transactions between owners, operators and third-party services within its system. As with DOVU, this vision aligns well with a fully scaled up seamless DLT implementation of MaaS. In the same vein, TSio Protocol (TTio Protocol, 2017) has published a whitepaper that describes how it seeks to build a fully integrated MaaS system based on blockchain. At its centre is a token that is tethered to an account, a geographic position and a device identifier thus enabling the system to track trips and the use of different services and vehicles and facilitate a single consumer-facing application while being able to distribute trip fees according to usage in the background.

The Mobility Open Blockchain Initiative (MOBI) is a consortium of car manufacturers, mobility, insurance, tech and energy companies, start-ups, NGOs, government agencies and academic institutions exploring the use of blockchain in transport launched in May, 2018 (MOBI, 2018). Transport-sector members include BMW, Bosch, Ford, General Motors and Groupe Renault. MOBI focuses on using blockchains in mobility services, auto manufacturing, vehicle data, cybersecurity, and tokenising related ecosystem transactions. MOBI will develop open-source blockchain software tools and standards to stimulate more rapid and scalable adoption of the technology by companies developing autonomous vehicle and mobility services While the consortia has plans to address a wide set of transport-related use cases, initial work targets the use of blockchains for securely establishing vehicle identity and history.

The Open Mobility System (OMOS) whitepaper sets out a comprehensive vision for a fully DLT supported application of seamless MaaS (OMOS, 2017). The OMOS model is different than those presented above in that it proposes a fully open and distributed governance model for its DLT-based MaaS ecosystem.

Numerous blockchain initiatives also target specific aspects of the mobility ecosystem. Porsche is collaborating with the start-up XAIN to trial applications for cars with a focus on blockchain-based use cases that are more convenient, faster or simply not possible with existing technology. These include temporary access authorisations (e.g. making it possible for an authorised person to unlock, but not start, the vehicle to deliver a parcel), remote locking or unlocking and secure and user-controlled data logging (Porsche, 2018).

A joint project by the Swiss Bank UBS, energy company Innogy SE and automotive technology company ZF demonstrated how blockchain technology could optimise charging patterns for electric cars by building blockchain-based eWallets for cars that would then negotiate charging requirements with energy suppliers based on pre-established and automatically executing smart contracts (ZF, 2017). Other electric vehicle charging applications include Share & Charge (Share & Charge, 2017) and the blockchain-enabled charging station Ethan BIoT (Blockchainfirst, 2017). Car or truck-based eWallets could also be used for toll applications such as proposed by Quantoz in its winning proof-of-concept entry at Kapsch TrafficCom (Kapsch TraffiCom, 2017).

More broadly, a number of blockchain and DLT-enabled regional and global exchanges for renewable energy are emerging that allow people to both produce and source energy, in even very small quantities, directly from each other in a trusted and secure manner without having to go through centralised distributors. This has implications for carbon reduction strategies in the transport sector especially as the share of electric vehicles is growing in many areas. Swytch (swytch, 2017) is one such application that has initially developed a proof-of-concept application in Seoul and five other South Korean cities to trade their rooftop solar and other renewable forms of energy for network-specific crypto-tokens. The exchange is set to scale up to provide a global platform allowing micro-producers of renewable energy to find a market that could lower the costs of renewable electricity.

Parkgene (Parkgene, 2017)is connecting drivers looking for a parking spot with owners of parking spaces. Upon reservation by the driver a smart contract on the Etherum blockchain is triggered through the Parkgene platform and the full parking fee is transferred into a temporary escrow wallet from which it is dispersed according to pre-established rules once the driver has vacated the parking spot with the majority of the fee distributed to the owner of the parking spot.

Vinchain (Vinchain, 2017) is addressing the issue of falsified information about vehicle status, from odometer mileage to accident history with a focus on the reselling market. It offers to create a global decentralized blockchain based vehicle information database with information collected and verified from original manufacturers and sellers, country registries, insurance and leasing companies and possibly through IoT applications like Slock.it that combine blockchain and communication technologies to track physical assets and their status in real time.

CarPass is another DLT-based vehicle telematics tracking application that builds on the concept of a "digital twin" for objects (Stöcker, 2017). In this case, the DLT is not blockchain-based as it uses the IOTA Tangle (described further on). The "digital twin" is cryptographically linked to an actual object in the real world (in this case a car) and stores that object's history (for example, "where did the object travel", "under what service agreement (for shared services)", "how was it maintained", "what is its insurance history", etc.) that then can be accessed by those with appropriate rights to maintain, charge, transfer ownership or any other operation that would require access to that history. The concept of DLT-enabled "digital twinning" also has many other applications in the delivery of MaaS (e.g. for shared bicycles, scooters, private cars, etc.).

Ensuring valid vehicle insurance coverage is another task that is well suited for DLTs. AXA Mexico has partnered with start-up Kayum under a Mexican Association of Insurers initiative to develop a DLT-based insurance validation blockchain protocol. The ease of use of the protocol and the trusted and immutable nature of the validation process should help ensure higher rates of insurance compliance in Mexico (only 30% in 2017) than at present (Etherisc, 2017).

Many governments are also exploring uses for blockchain and other DLTs for regulatory tasks. The governments of the Netherlands, the UK, Dubai and Taipei have announced that they will explore, introduce and support blockchain or DLT initiatives in a number of areas, including in support of transport and smart city initiatives. Estonia has been a leader in the field of e-government with a broad vision that was first put into place in 1997 (e-estonia, 2018). Since 2008, DLT based on cryptographic hashing has been a core foundational technology for the e-state as a way to securely manage sensitive personal and commercial identity. E-identity is the immutable identifier that allows citizens to interact with, and control access to, their data when dealing with tax authorities, healthcare providers and other e-service domains. In the case of transport, the DLT-based identifier has enabled authorities to find a solution to an otherwise challenging problem – how to quickly and reliably ensure that ride service drivers report and pay taxes. The existence of the DLT-powered e-identity made it possible for the Estonian Tax and Customs Board to negotiate automatic tax filing and payment protocols with both Taxify and Uber. These arrangements allow for the automatic payment of ride service-based taxes and have led to a 460% increase in the number of declared drivers and a 660% increase in declared income from ride services (Estonian Tax and Customs, 2017).

Applications for freight transport

Freight transport and supply chain optimisation and improved, authenticated clearing functions are another area where blockchain and other DLT technology holds promise. DLTs offer a trusted and quick authentication method for goods and their status along the whole transport chain, including verification of driving distances and adherence to driving laws and customs regulation. In this way, blockchain and other DLTs offer a way to address fraud, theft and systematic inefficiencies that otherwise drive up the cost of shipping and logistics chain management.

Various players in the sector have either initiated their own DLT trials or joined an alliance, e.g. the Blockchain in Transport Alliance (formerly known as the Blockchain in Trucking Alliance) that positions itself as a forum for the development of blockchain standards and education for the freight industry. This alliance recognises that blockchain technology could not only simplify procedures within the freight transport sector but across the whole transport sector. Its partners are global and among them are UPS, FedEx, Penske Logistics, GE Transport, SAP, Daimler, etc.

Individual companies have also been active in deploying DLT applications for freight transport. For example, in early 2018 Maersk, a global container logistics company and IBM formed a joint venture to develop a global trade digitization platform built on open standards and using blockchain technology (Hackett, 2018). It focuses on establishing two components: first a shipping information pipeline that provides permissioned end-to-end supply chain visibility for all actors involved and allows them to manage their operations securely and in real time. Secondly to digitise and automate paperwork that can be securely submitted,

validated and approved across organisational and national boundaries. Crucially, the joint venture collaborates with customs and government authorities, Singapore Customs, Peruvian Customs and the Guangdong Inspection and Quarantine Bureau in the People's Republic of China (PRC) among them, to enhance supply chain security and facilitate trade flows.

Blockchain, distributed ledgers, Bitcoin and (other) cryptocurrencies: What's what?

Much of the attention generated around blockchain applications has been in the context of Bitcoin and other cryptocurrencies. These have garnered considerable, likely over-hyped, attention. While cryptocurrencies have implications for the transport sector and for MaaS applications (storing value, payment, taxation, user fees, etc.), what they offer is different and separate from the core functionality and promise of DLTs for transport.

The suitability of DLTs for early stage cryptocurrency applications is relevant, but not critical to the potential of DLTs for longer-term transport applications – especially if accompanied by changes in the regulatory treatment of DLT-mediated transactions. This poses a challenge for policy since the long-term potential of DLTs is not clearly discernible from early and limited use cases – especially as they are still co-evolving with other broader societal and technological changes.

There are three broad generations of blockchain/DLT implementations (Pavel, 2017):

- DLT 1.0: All that relates directly to the creation, transfer and payment functions of cryptocurrencies
- DLT 2.0: All of the financial and economic applications building on Blockchain 1.0 (Shares, term contracts, loans, intellectual property, smart contracts, etc.)
- DLT 3.0: All other applications of DLTs outside of the economic and financial areas – including IoT (and therefore MaaS), healthcare, administrative functions, asset tracking, science, art, etc.)

Much of the current body of knowledge and experience in DLTs relates to generations 1.0 and 2.0 but most of the applicability of DLTs for MaaS lies in generation 3.0 which is still very much under development.

Moving beyond centralised ledgers

Ledgers are at the core of almost all economic transactions. They record events that occur in the real world ("s has paid y to z in return for good/service a", "vehicle a is at location b and has capacity c", "traffic at time t and location y was v") and associate these events with a time-stamp that establishes the authentic and immutable sequencing of events ("event a happened before event b and therefore outcome c can proceed").

A core functionality of ledgers is to associate events with people or other entities like businesses in an agreed version of reality. Ledgers are kept by those that have a stake in ensuring the exactitude and veracity of their contents. Ledgers have historically been used most commonly to track fungible assets like money or other asset types like property. But ledgers can encompass any time-stamped event linked to a broad class of assets. For example, a public transport company keeps a ledger of who has paid what for subscription to their services, a ride service company keeps a record of who has used its services, where they have travelled and how much was paid, tax authorities keep track of what is owed by entities and how much of this has been paid, a car manufacturer keeps track of vehicle serial numbers and information related to the construction and disposition of vehicles, an insurance company keeps track of who has paid for what coverage, a shared bicycle provider keeps records of how it's bicycles have been used and by whom.

Ledgers and their contents underpin many business and regulatory processes and have considerable value for those who control them. These agents spend considerable resources ensuring that what is contained in

their managed ledgers is trustworthy and therefore can be used by different parties at face-value. The management of traditional ledgers also typically involves large-scale, oftentimes bespoke, legacy computer systems and data formats. The inherent value recorded in ledgers explains why they are typically developed and managed in closed silos. This allows for centralised control in the management of transactional information, identity and rights. But this strength can also be a weakness in that it opens up the possibility of single-point failure. This vulnerability is countered by internal redundancy and backups of ledgers – but these just reduce rather than eliminate the risk of single point failures.

Cloud computing created new opportunities for securely handling ledgers and online databases. Rather than distributing information within physically siloed databases, cloud-based databases centralise these functions in a virtual database housed and managed seamlessly across multiple distributed servers. Rules regarding data access and sharing are determined by internal protocols and permissions set by the cloud provider and agreed by all users. The cloud is managed by the cloud manager/owner. Though cloud-based ledgers enhance redundancy and allow for robust access-control because of centralised functions, they are still largely centralised ledgers and are vulnerable to many of the same risks.

Underlying principles of distributed ledger technology

There are many types of distributed ledger technologies and many iterations of their constituent protocols. Nonetheless these share many core features. The case of blockchain helps illustrate how these distributed ledgers function. Blockchains have seven principal characteristics:

- Distributed databases and ledgers;
- Irreversibility of records;
- Transparent identity management with pseudonymity;
- Robust validation and consensus;
- Peer-to-peer transmission;
- Computational logic;
- Distributed databases and ledgers.

Distributed ledger technologies move away from the core logic of traditional ledger frameworks. Unlike traditional ledgers, DLTs are not siloed but, rather, identical copies distributed across all users. Each party on a blockchain has access to the entire database and its complete history. No single party controls the data or the information. Every party can verify the records of its transaction partners directly, without an intermediary. In some cases, certain data linked to the transaction record may be only accessible to entities with sufficient access rights as set by the protocol and verified by digital signatures. Blockchains can be permissioned – in which only vetted parties ("nodes") can edit and add to the blockchain (e.g. in a closed network), or open, in which any party engaging in an approved transaction can add to the blockchain (Figure 8). Block validation (see discussion further on) is simpler and more straightforward in permissioned vs. un-permissioned ledgers.

While cloud-based ledgers allow multiple agents to access a central database, DLTs distribute the copies of the database itself so that all permissioned users (or any user in permissionless open system) can cross-check its exactness and the veracity of its contents at any given time.

Figure 8: **Taxonomy of ledgers**

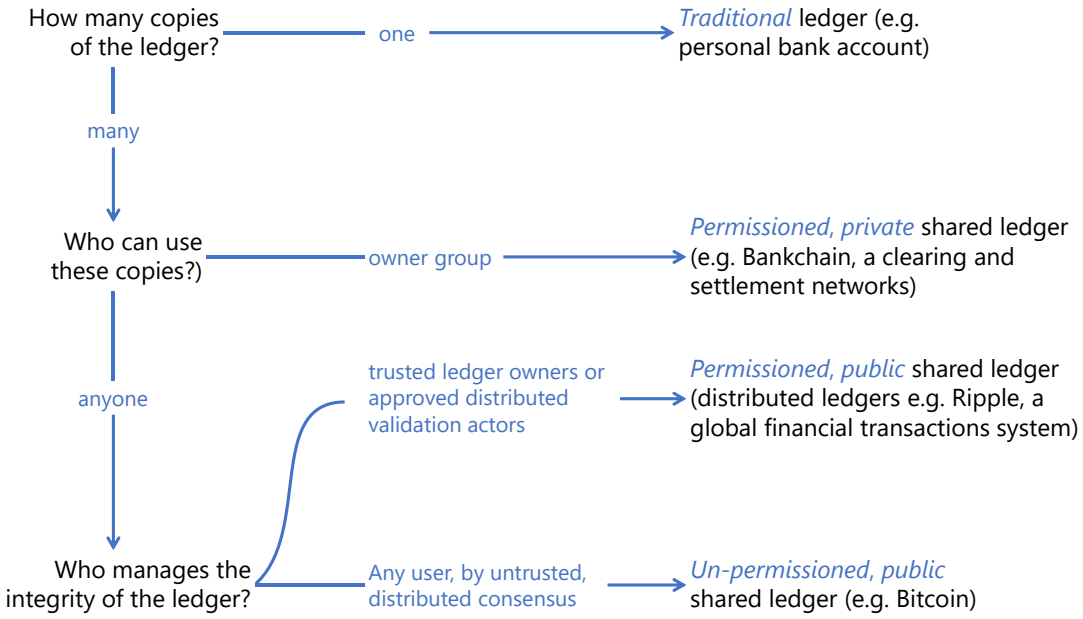

Source: **(UK Government Office for Science, 2016)**

Irreversibility of records

Once a transaction is entered in the database and the accounts are updated, the records cannot be altered, because they're linked to every transaction record that came before them (hence the term "chain"). The recording of the blockchain database at any given time is permanent, chronologically ordered, and available to all others on the network. This immutability is at the heart of the "trustfulness" of the blockchain.

Figure 9: **Hash-linked blocks in a blockchain**

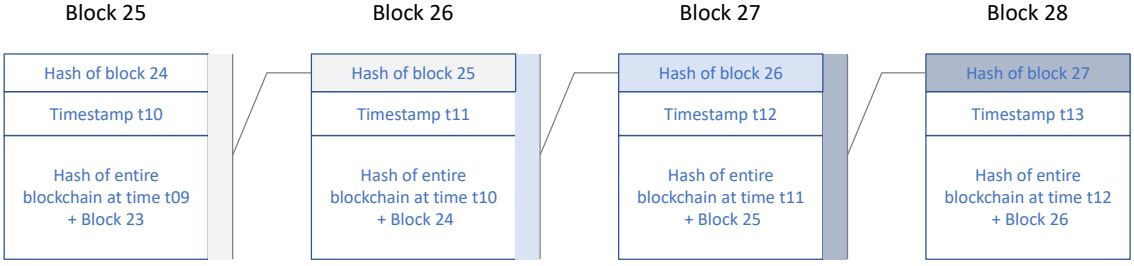

Each event recorded in a blockchain gives rise to a new "block" that is identified by a unique cryptographic identifier (Figure 9). The cryptographic identifier in each block is linked to the preceding block's identifier (blocks are "chained" together). Each cryptographic identifier references not just the heading of the current and preceding block but all of the data within the preceding block (and thus all of the data going back to the original block – the "genesis" block).

These cryptographic identifiers are created on the basis of "hashing" or algorithmically transforming an arbitrary amount of input data into a fixed-size output – the "hash" (Figure 10). This means that a two-character input will give the same size hashed output as a 10 000-character input. There are numerous available hashing algorithms – the one used by Bitcoin is SHA-256 which returns a fixed 256-bit output from any given input. Though the entire blockchain of Bitcoin is over 157GB in size as of early 2018, all

participants need to cross-check the validity of the copies of the blockchain they hold is only the 256-bit hashed header of the most recent block. Hash size-consistency is an essential feature of the blockchain and allows for rapid processing of blocks irrespective of the size of the content they reference.

Figure 10: **"Hashing" characteristics in support of distributed ledger technologies**

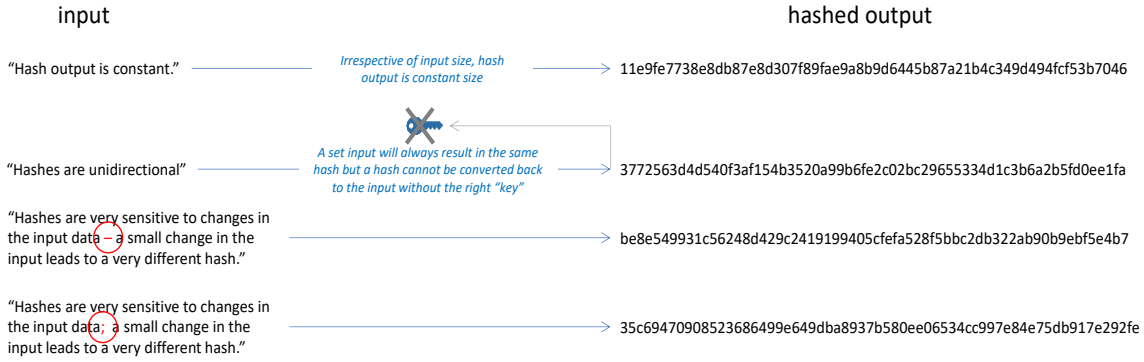

The immutability of blockchain stems from its state-dependency. Hash output is very sensitive to even very small changes in the input data – any change in the input leads to a very different hashed output. This means that any tampering with the content of a block (or the hash of the content) or its header would lead to a change in its hashed identity and would lead to a new hash header for that block. This would no longer match all the other distributed copies of the blockchain and would be flagged as an inauthentic copy.

The security of a blockchain is a result of its unidirectional nature (Figure 10). A hash can be generated from any input data, but the input data cannot be elucidated or reverse-engineered from the hash. In addition, security protocols might also call for double-hashing (as with Bitcoin). An initial hash output can itself be re-hashed before its use as a pointer in the blockchain.

Non-discovery via hash uni-directionality and irreversibility are at the heart of the DLT/blockchain model. These features are "hard-baked" into the protocol and, alongside the distributed nature of the ledger, they ensure the "trustability" that any event described in any given block at the time of its timestamp was an accurate and immutable description of reality without having to reference a central authority or database controller. The importance of robust hashing underscores the importance of using well-tested and accepted cryptographic encryption protocols in support of DLTs. These may evolve over time and so, ideally, DLTs should be built to account for changes in cryptographic security or have contingencies for ledgers (or parts of ledgers) that are may be vulnerable due to evolving cryptographic capabilities.

Transparent identity management with pseudonymity

Identity and authentication are ensured using digital signatures. Digital signatures are a way of establishing irrefutable identity cryptographically. They are used to link block creation to a single, authenticated entity and can be linked to rights to view or access the encrypted contents of a block, of its hash or any other encryption-protected data object. In cryptocurrency applications, digital signatures authenticate ownership of assets and the rights to dispose of them.

Asymmetric encryption algorithms generate a pair of keys comprised of a "public" key that is distributed and "private", non-shared key used to digitally sign data objects. Each node, or user, on a blockchain has a both a public key that identifies it and a private key that serves to authenticate its identity via digital signatures. The asymmetry in the encryption stems from the fact that it is always easy to elucidate a public key from its private key but essentially impossible to elucidate a private key from its public counterpart. Every transaction and its associated public key are visible to anyone with access to the system. The two keys are mathematically linked to each other so that if a data object digitally signed by the holder of a

private key does not correspond to that entity's public key, then that data object or transaction is deemed inauthentic.

The private key holder has rights associated to blocks that they have created or to which they have been granted access via the block's private key-holder. This means that it is possible to assign rights to access data within a block or to build identity-conditional rights into code that may be activated within a block.

Public-private keys and digital signatures, especially when combined with computational logic and smart contracts can be a strong enabler of everything-to-everything MaaS. For example, user A and user B both travel using the same service. Their travel activity is recorded in blocks that are created and digitally signed by each. User A wishes to use another service (not necessarily a transport service) that requires limited access to her travel data. She allows the other service provider to access part of her encrypted data as authenticated by her digital signature. User B does not wish to provide access to his data and thus does not gain access to the other service. In the end, the new service provider only has access to user A's data for the limited set of uses user A has agreed.

Transactions occur between blockchain public key addresses, thus while these transactions and their public keys are transparent and visible to all, they are still pseudonymised. Again, cryptographic methods are central to the way in which DLTs manage identity and access rights and so the use of well-known and robust cryptographic protocols is essential.

Robust validation and consensus

Individuals, institutions and governments routinely interact and share information about transactions and events in which they are involved. All parties must trust that information that is shared is authentic, that all parties are who they say they are and that what has happened has indeed occurred. This trust is delivered through vetted identity management and data recording in centralised databases and ledgers held by trusted authorities. As noted before, this model has strengths but also some inherent vulnerabilities as well.

One of blockchain's and other DLTs' fundamental innovation is that they are systems set up to bypass reliance on any centralised institution or reconciliation ledger. The way DLTs are set up replaces trust between different parties or trust in some form of an oversight committee with cryptographic proof of validity or "consensus". It proposes that any transaction could be authenticated and any transmitted piece of information maintained by an emergent process of consensus among a globally distributed network of peers that follow a precise, incorruptible method to check any change in the system.

The cryptographic identity of each new block in a blockchain must be validated (recognised as authentic, in conformity with the cumulative blockchain and linked to a unique identity) before it can be included in the latest iteration of the ledger that is propagated to, and recognised by, all nodes. These "consensus" protocols and algorithms are linked to the nature of the blockchain in question.

Bitcoin (open and permissionless) uses "proof of work" (POW) to validate new blocks. This consensus protocol is costly in terms of computing power (and energy – see Box 1) since its difficulty scales with the size of the blockchain. Because of this, validating nodes ("miners") are incentivised to run these computations by receiving some stake in the value of the blockchain – 12.5 Bitcoins per block validated in February 2018. When a new event occurs and is to be recorded (concatenated to the current global copy of the ledger) it is bundled with other transactions in a block and distributed to the network. Any node in the network can compete to be the first to "mine" the right solution to the cryptographic puzzle that would ensure that the contents of the proposed block are legitimate (i.e. that they reflect the unadulterated sequence of transactions in the blockchain).

The Bitcoin validation task, like many blockchain validation processes, requires miners to find an unknown integer (the "nonce"). This integer is concatenated to the contents of the block awaiting validation, the digital signatures of the parties and the hash of the previous blocks in the chain of transactions. This

concatenated text-string is then parsed through the SHA-256 hashing algorithm and results in a hashed output that must meet specific conditions – namely that the hash starts with a specified number of zeros. The required number of zeros is determined by the current level of "difficulty" set in the Bitcoin blockchain protocol and increases as the size of the blockchain grows. Finding the correct "nonce" under these set of conditions requires randomly iterating integers and passing them to the hash function as per the protocol until one results in the specified outcome. This requires time- and computing resource-intensive brute force calculations which only become more difficult as each successive block is validated or "mined".

Box 1: **Distributed ledgers and energy consumption**

The energy consumption of blockchain technologies is a direct outcome of the trust-building mechanism DLTs use to ensure security and pseudonymity: the so called "proof of work" (POW).

POW is intentionally set up to be mathematically difficult which increases the processing power and time required to solve it. This results in elevated energy consumption by the machines processing the POW. Currently the POW is not only used by the Bitcoin blockchain, but also by Ethereum and many other DLTs, even if some specific functions of the underlying hashing procedure differ amongst DLTs. The difficulty to perform the POW is a critical factor in building the trust in the dispersed blockchain network.

Next to computing hardware, electricity is one of the biggest cost factors for performing POW calculations. POW processors (either in general-purpose computers or dedicated mining rigs known as Application Specific Integrated Circuits – ASICs) need electricity to perform the POW calculations and also to cool themselves down. The server farms require dedicated buildings and some supervisory staff as well – including security guards to prevent theft of ASICs. For these reasons, many mining centres are set up in regions that are naturally cool and also have a lot of inexpensive electricity and land, e.g. Iceland with its cold climate and large geothermal energy supply or Sichuan province in China where land is inexpensive and hydro-electric power plentiful and cheap.

The more a specific blockchain is adopted, the more attractive it becomes to be the first to perform the POW and reap the connected reward, e.g. a certain amount of Bitcoins. Thus the more a specific coin is valued, the more attractive it becomes to spend even more on fast computing power and with it more on energy. At the same time, the "difficulty" of the POW calculation is set to increase in line with the size of the blockchain and the hashpower deployed. This means that greater energy efficiency in Bitcoin mining machinery does not lead to a reduction of overall energy use – on the contrary.

Calculating overall blockchain-related energy use is not straightforward. Despite different statistics that have been published likening the energy consumption of the Bitcoin network to the energy consumption of Denmark or Cyprus, or to a full 1% of the overall energy consumption of the USA, getting exact energy consumption figures is not easy and the estimations rely on a broad range on assumptions that are difficult to verify. One bottom-up analysis that seeks to identify, locate and calculate the energy consumption of mining facilities carried by the University of Cambridge Centre for Alternative Finance (CCAF) estimates that at least 462 megawatt hours (MWh) (0.000462 terawatt hours - TWh) are required to secure the Bitcoin blockchain alone (Hileman & Rauchs, 2017). In contrast, one of the most cited Bitcoin energy consumption estimates carried out by Digiconomist estimates Bitcoin's current energy consumption to be 30.2 terawatt-hours (TWh). This is orders of magnitude greater than Cambridge's estimate and represents more energy than 63 different countries consume annually (and roughly equal to the annual energy consumption of Greece).

These estimates are several thousand orders of magnitude apart and highlight the need to better estimate and monitor energy use in relation to blockchain POW. As described in the text, there are numerous ways in which the energy cost of POW and validation can be lightened, but for now, it seems clear that the POW methods currently implemented are unsustainable when it comes to energy.

Once validated, the solution, including the "nonce" is transmitted to all the nodes in the network. While finding the solution (and therefore validating a block) is difficult, verifying a solution (and therefore ensuring that the block is "authentic") is easy since all a node must do is concatenate the correct "nonce" to the other elements of the block and ensure that the resulting hash starts with the right number of zeros as specified by the protocol. In the case that two competing miners find the same solution, the Bitcoin protocol stipulates that the longer blockchain is considered to be the valid one.

This process is relatively slow and leads to a high latency in consensus formation which is prejudicial to the usefulness of the Bitcoin protocol for scaled-up and low-latency uses. This is especially the case for certain potential applications of DLTs in the MaaS ecosystem or, more broadly, for applications in support of "internet-of-things" (IoT) applications.

Alternative consensus protocols, essentially in the domain of cryptocurrencies, have sought to reduce latency and therefore may open the pathway for more widespread use of DLTs for transport broadly, and for MaaS specifically.

One approach, "proof of stake" (POS), built into the Ethereum (another cryptocurrency) protocol allocates validation tasks to a smaller set of nodes based on the stake they hold in the blockchain. These nodes place some of the value they control in escrow and gain proportionate validation responsibilities and rights under the assumption that nodes with the highest stake in the blockchain will be those most incentivised to keep it accurate.

Another consensus protocol, "Proof of authority" (POA), is especially suited for closed and/or permissioned DLTs. Instead of validation authority being allocated according to the stake, or "skin in the game" held by a node, validation authority is allocated to selected nodes based on their actual authenticated "real-world" identity. For this to work, there must be a robust and standardised way to ascertain identity and ensure that potential validators are indeed who they claim to be. The right to validate should be scarce – earning it should be difficult, retaining it should be valued and losing should be unpleasant (reference POA source). Such an approach improves the speed of validation but comes with vulnerabilities inherent in having a small set of validating nodes and authorities and thus the risk of single-point-of failure attacks or collusive behaviour among nodes.

Peer-to-peer transmission

Communication of the blockchain occurs directly between peers rather than through a central authority. Each node stores and forwards information to all other nodes. Peer-to-peer transmission of the entire database is not frictionless and is conditioned by the speed with which data can be propagated to all participating nodes in a blockchain. This speed is a function of internet data Transmission Control Protocols/Internet Protocols (TCP/IP), and of data transmission speed across communication networks. Latency in the propagation of copies of a blockchain may lead to a situation where some participating nodes have a more current version than others. This leads to a "fork" in the blockchain where records of a same blockchain no longer correspond to each other since some have newer entries. These differences subside relatively quickly as validation protocols confirm the authentic sequencing of block creation unless, as discussed later, a specific brute-force attack tries to prolong the inauthentic chain.

Computational logic and smart contracts

The digital nature of the DLTs means that recorded events and transactions can be tied to computational logic and in essence programmed. This means that users can set up algorithmic rules (e.g. "smart contracts") that automatically trigger transactions between nodes when the right set of conditions are met. Smart contracts were originally included in the Ethereum blockchain protocol and have since been built into other DLTs.

If a blockchain or other DLT is the ledger that records information, the smart contract is the operational application layer that initiates transactions. Smart contracts can govern specific business processes (e.g. if user A has paid for service S, then user A can access that service, otherwise, access is refused) or any other set of "if-then" conditions – including enacting regulatory oversight or control (e.g. if driver B is licensed for activity A and is registered with tax authorities for automatic declaration, than driver B can sign on to platform P to carry out activity A).

Traditional contracts are typically drawn up in isolation and irrespective of the computer code that might execute them. The contract is drawn up in human language, recorded on paper or its electronic equivalent, physically signed and archived in such a way that it can be retrieved in physical form if needed in dispute settlement. If the contract terms are to be enacted via computer code, the latter is typically a non-standardised, legally unbinding approximation of the original paper contract which retains primacy for any settlement resolution.

This framework has been effective in governing straightforward, limited-party and static relationships but it is less adapted to govern large, multi-party, dynamic relationships such as those that are necessary to deliver complex services to a single user. This is precisely the challenge of delivering MaaS to users in a mesh-y world of government agencies, service providers and technology.

Further, smart contracts can help realise significant efficiency improvements in business and regulatory processes by leveraging robust identity and authentication mechanisms built into DLTs to reduce unnecessary duplication of data recording and retrieval (with the inherent risk for error that these entail). In this world, data entered once is available to all business and regulatory processes according to rights assigned by the data subject or owner. Smart contracts can also build on the potential for public/private key parametric privacy to create new value for users of services without compromising their privacy as discussed earlier.

Open technology questions and challenges

For all of the promise that DLTs may hold for widespread and decentralised MaaS applications, they also have some inherent vulnerabilities. These relate to the three principal and interlinked areas:

- **Scalability**: Are DLTs like blockchain able to scale up and handle much larger and widespread tasks than cryptocurrency (but as well for cryptocurrency applications)?
- **Speed**: Will DLTs be able to mediate transactions and events with extremely low latency (and even in real time) as required by large-scale MaaS applications?
- **Security**: Are DLTs' core foundational protocols, especially as they relate to cryptographic technologies, vulnerable to their design limitations or robust to developments in computer science, and in particular to quantum computing?

Scalability

If DLTs like blockchain cannot scale to potential MaaS use cases involving millions or billions of transactions, then they will have limited applications, especially for creating an "instant" marketplace and transaction space bringing together multiple operators, authorities, asset owners and users.

The size of an open and permissionless DLT like Bitcoin is one factor to consider in the scalability challenge. Large file sizes may present storage and bandwidth challenges that limit the scalability of the blockchain. This challenge, however, may not be as critical as some other factors since storage capacity and bandwidth transmission rates have also historically scaled up.

There are multiple strategies that have or are being deployed to address this issue. Protocol-level storage fixes targeting nodes may include storing partial reference copies of the full ledger (still tethered to full copies held by certain nodes) e.g. "pruning nodes". Another strategy is to create "side-chains" that are pegged to, but operate off of (and later reconnect to the main blockchain -- e.g. Raiden's technology for Ethereum chains (Raiden Network, 2017). Ethereum supports creating state-dependent blockchains -- e.g. "sharding". Shards relate to one asset only, or one transaction primitive (Ethereum, 2018) and can be linked as necessary to the broader underlying chain. Finally, another promising strategy is to move most of the space-consuming transaction data off of the blockchain while retaining the hash-ensured trustability and security features of the full blockchain (Lightening Network, 2018). Many of these approaches have the

added benefit that they allow certain classes of transactions (usually small or low-value transactions) to be validated offline and reconciled with the core chain when the transacting entity is back online.

Box 2: **SAS Event Stream Processing: Blockchain application testing and analytics**

As the use of blockchain and other DLT increases, transport organisations will need to access and analyse data from different blockchain processes and applications. These analytical needs pertain to two types of data:

- Data at rest – static data that already exists in a blockchain's immutable data store
- Data in motion – data that is being produced every time a transaction is created in the blockchain.

Exporting static blockchain data into an analytics platform provides insight into various transaction characteristics, including: trends, segmenting transactions, predicting future events, and identifying relationships between blockchain and other data sources. With the advent of streaming analytics, blockchain data in motion offers additional opportunities for analysis, which can help organisations, especially government agencies, identify changes in near real-time in blockchain-based processes. Seeing these changes as they're happening provides an opportunity to take immediate action to address these changes.

Analytic models developed using static data can be deployed in real-time to ensure the integrity and authenticity of blockchain transactions. A good example is identifying and combatting real-time payment fraud in-transit. Fraud in transit payment systems (especially in heavily used regions such as European countries) is common. Blockchain analysis in real time can identify the fraudulent activities and deny any suspicious transaction in advance. But a key challenge is how to analyse streaming data.

One approach - Event Stream Processing (ESP) pioneered by SAS centres on a blockchain simulator to demonstrate the application of real-time blockchain analytics. This simulator generates client requests into a miner process that is controlled by a consensus validation process. Both the simulator and consensus processes use the publish/subscribe APIs connected to the model for managing blockchain updates.

Figure 11: **SAS Event Stream Processing blockchain simulator**

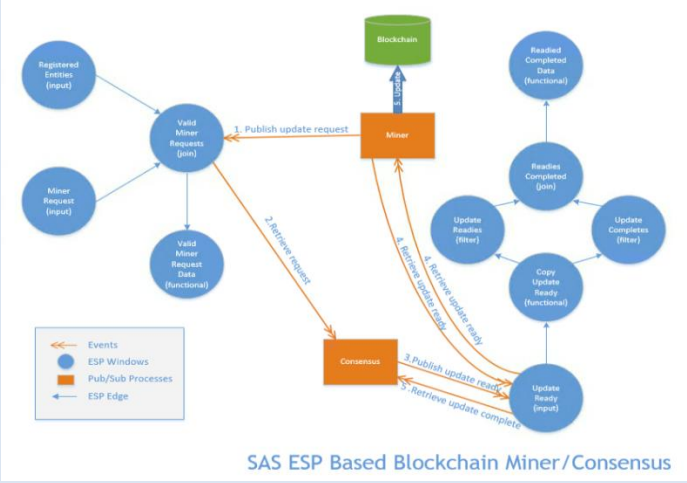

This approach produces operational streaming analytics covering transactions per second, block updates per second, and total transaction times from creation to block update. The processing parameters can be changed on the fly via a parametric user interface. Deep learning algorithms at the miner and consensus process levels automatically manage blockchain metrics such as block size and elapsed time. Event Stream Processing can quickly help assess scaled-up Internet of things-type applications. As blockchain technologies mature and IoT use cases become the bellwether for blockchain implementations, the need for higher speed block updates, processes and communications will trend toward stream-based composition. This, in term, means that analytical capacity is able to handle these applications, as demonstrated by the ESP simulator.

Source: https://blogs.sas.com/content/sascom/2017/12/15/practical-approach-blockchain-analytics/

The scalability challenge of DLTs, and especially of Blockchain, is also linked to the architecture of various DLT protocols. Bitcoin blocks, for example, are limited to 1 Mb in size and Ethereum has a *de facto* scalable cap based on processing effort and bandwidth/storage constraints that miners are willing to accept and pay for. These limits in turn have an impact on block creation and validation rates. As of February 2018, the 60-day running average Bitcoin block transaction confirmation time was 11 hours (Blockchain.com, 2018). These size limits may also have an incidence on node synchronisation times, especially for new nodes. One much-discussed approach is simply to increase Bitcoin block size. This would allow for much greater scalability but would also increase the amount of computational power necessary to mine and process blocks which would de facto lead to a concentration of validation power to those entities best able to mobilise and deploy the considerable resources necessary. This concentration may already be seen in the case of Bitcoin as discussed further on.

Ledger size bloat is likely for DLT applications that contain "logging" data as opposed to simple transaction data. This raises the question of how much can be pushed off the distributed ledger and just cryptographically referenced by the latter using permissioned access rights built into blocks. This is essential in the case of IOT applications and raises the question of offline validation. Real-time analytics of blockchain and other DLTs is an emerging area where solutions are just now being deployed (Box 5).

Speed

If DLTs cannot handle a very large number of events in a very small amount of time, they will likely not scale beyond a few high-latency tasks despite their other features. This is especially true if they are to serve as the basis for delivering seamless and distributed MaaS services. The speed of DLTs, as for size, is structurally linked to protocol definition (especially validation and consensus mechanisms), which, in turn, is often linked to a specific security concern. The Bitcoin protocol, for example, specifies that average block creation times (e.g. transaction times) should be about 10 minutes. Ethereum block creation times average about one every 14 seconds (Etherscan, 2018). These limits are built-in security features meant to limit blockchain vulnerabilities to nefarious use – e.g. countering the risk double-spending due to too-fast creation and propagation of identical blocks.

Current blockchain transaction rates (measured in transactions per second or Tps) are very low – about 7 Tps (Figure 12). Ethereum transaction speeds are faster, at about 20 Tps, but still orders of magnitude slower than those achieved by the Visa network or even PayPal. Sharding and other protocol changes show potential to match or surpass Visa rates but these have not yet been realised at scale. Many of the mechanisms discussed (increasing block size, off-chaining, sharding, etc.) promise to increase transaction speeds and the usefulness of blockchain-applications in support of low-latency MaaS applications.

Figure 12: **Transaction speeds for payment services and blockchain cryptocurrencies**

Source: (Amoros, 2018)

A promising longer-term approach to the speed/scalability challenge of DLTs is to move away from the blockchain model to an entirely different distributed ledger model. One that is purpose-built for speed but

that still leverages the distributed, secure, trusted and pseudynomised qualities inherent to blockchains. Directed Acyclical Graph (DAG) technology is such an approach that seems well-suited for large-scale IoT and MaaS applications because they are scalable and fast due to their structure and consensus mechanisms.

Like blockchains, DAG-based DLTs such as ByteBall, Hashgraph Swirlds and IOTA link transactions via cryptographic keys and hashed identifiers. Unlike, blockchain, however, DAGs involves no mining and linking of blocks in linear transactional chains. Rather, each transaction is linked to a small group of previous transactions in an emerging lattice-like chain. Validation and consensus models are different as well in that each new transaction must validate a constrained set of past transactions to which they are linked and not the whole transaction chain. The protocols for doing this vary but are built on the same approach.

IOTA (IOTA, 2018) is an indicative case of both the potential for DAG DLT implementations for internet of things (IOT) tasks like MaaS and illustrative of some of the pitfalls that are related to early implementations of any DLT.

IOTA has been developed specifically for IOT applications involving large, heterogeneous, networks of transacting entities (e.g. sensors or vehicles). IOTA is built on a DAG structure called the "Tangle". The Tangle is a blockless DLT that builds consensus directly into its architecture. Each new transaction added to the tangle must validate two previous transactions. Transaction recording and validation in IOTA is a three-step process.

In the first step, the new transaction is created and signed cryptographically. In the second step, the new transaction (a "tip") is cryptographically linked and bundled with two other previous transactions. In the third step, the IOTA protocol defines a nonce-based "proof of work" outcome that must be met for each of the two past transactions in the bundle. In scenario A (Figure 13), Tip T_1 confirms transactions 26 and 27 and Tip T_2 confirms transaction 27 and re-confirms transaction 26. The Tangle protocol is engineered such that Tips validate a mix of new and recent transactions thus building consensus. As more tips confirm a same transaction, confidence in the validity of the transaction grows. The Tangle proof of work is similar to, but more lightweight than the POW constraint set in Bitcoin. Tangle POW helps to ensure that transactions do not expand invalid extensions of the graph. Once the "proof of work" condition is met, the Tip is considered accepted and the transaction is confirmed.

The IOTA Tangle protocol allows for rapid transaction confirmation and, unlike blockchain protocols, validation speeds scale with the size and complexity of the tangle. Another feature of IOTA's Tangle DLT is that the model calls for fee-less transaction recording. Whereas blockchains like Bitcoin must incentivise "miners" to carry out resource-heavy POW-based block validation (and thus transactions incur a fee), IOTA's Tangle model requires no mining, no incentivising of miners and no fees. IOTA's feeless model is especially interesting in IOT deployments like MaaS where micro-transactions between connected objects, vehicles and parties would be the norm. In a blockchain model, fees may outweigh the value of micro-transactions and that could limit the scope of potential MaaS transactions.

Another potential use of DAGs like IOTA beyond logging data from multiple sensor platforms for the delivery of services is the collection of data for testing purposes. The International Transportation Innovation Center (ITIC) has announced that it will use IOTA's Tangle DAG to build a global network of testbeds to develop and validate technologies that support connected, automated and zero-emission driving (ITIC, 2018) This would enable a wide range of automotive manufacturers and service providers to make available vehicle testing data so that all can benefit from safety-relevant data.

Figure 13: **IOTA Tangle Directed Acyclical Graph-based distributed ledger**

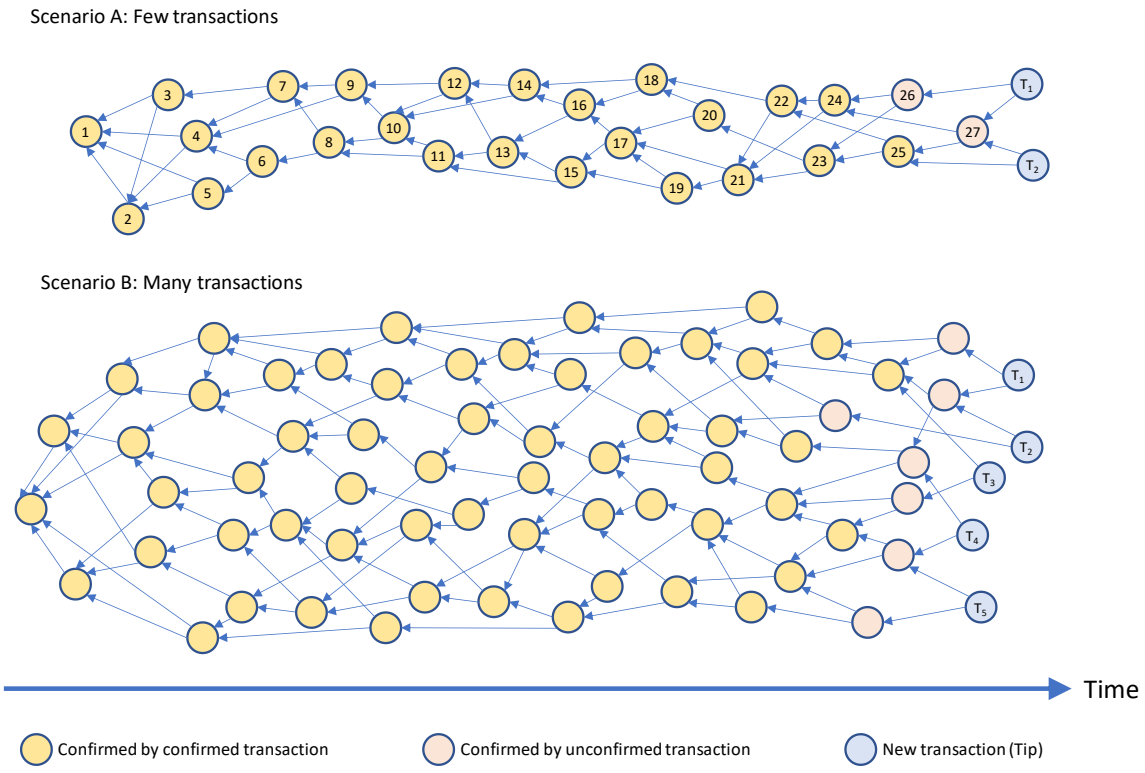

Source: Based on IOTA, 2017

Speed, scalability and suitability to micro-transactions all make the IOTA DLT model an interesting one for MaaS. But just as with everything else with DLTs, while the proof of concept may be enticing, the ultimate applicability of DAGs like IOTA to real-world applications is uncertain and evolving and will depend on more than the theoretical applicability of the technology to a concrete application, like MaaS.

For example, the initial implementation of IOTA incorporated serious vulnerabilities built into its bespoke hashing algorithm ("Curl") (Narula, 2017). These vulnerabilities have since been addressed and patched but researchers have pointed out other issues as well. IOTA also uses ternary vs binary number encoding which in itself isn't a vulnerability but it limits IOTA's compatibility with "off-the-shelf" algorithms and security analysis software. A well, IOTA's early-deployment dependence on a "trusted coordinator" function to protect against hijacking the transaction validation process has been seen by some as addressing a structural weakness in the protocol that may persist even at scale (IOTA is centralized) (The Transparency Compendium).

DAG-based DLTs like IOTA show how some of the current limitations of blockchain technology could be addressed, including those linked to speed and scalability in complex, multi-party and asset implementations. But it also shows that new DLT models must not create new vulnerabilities as they address limitations with early blockchain technologies.

More broadly, focussing on current DLT transaction speeds and benchmarking them against other networks like Visa's may not ultimately be a helpful approach. It is somewhat like comparing early electronic messaging speeds (2.4 kbit/s in the original ARPANET, the predecessor to the internet (htt)) to early facsimile (fax) transmission speeds (2.4-4.8 kbit/s (htt1)). Despite near-parity in these speeds, the

ultimate value of the internet has proven to be orders of magnitude greater than facsimile transmission technology.

Size, storage requirements and transmission speeds were also flagged as early constraints to the scaling up of the internet and yet these have developed as the value of a global network of interconnected computers running common data transmission protocols has grown.

The point is that if the principles behind DLTs are compelling enough, and they seem to be for certain use cases, then solutions are likely to be developed that build on their potential. This means that DLTs may "grow into" use cases such as MaaS rather than be directly deployed in their current form. Again coming back to the fax analogy, limiting the discussion of how well DLTs can contribute value by comparing them to other transaction-processing systems in operation today is somewhat akin to having a discussion in the early 1980s about the future potential of the internet by only discussing how well electronic messaging could handle the number of faxes in circulation then.

Security

Blockchains and other DLTs address certain security risks and vulnerabilities in current data storage and transaction processing protocols. In particular, they largely remove risks associated with centralised data management and transaction processing. This is an important characteristic since the increasing sophistication and frequency of cyber-attacks increase both the likelihood and potential damage from single-point-of-failure exploits. Nonetheless, blockchains and other DLTs also display a range of vulnerabilities that must be understood in order to assess their potential, especially when compared to legacy systems that may have decades of cyber-protective engineering built in.

Most of the security risk assessment of DLTs concerns cryptocurrency implementations of blockchain technology. This is not to say that other forms of DLTs are less risky than blockchains (as discussed above) but simply that non-blockchain DLTs are much newer and have undergone more limited risk analysis.

The broad taxonomy of blockchain security risks can be broken down between those that are inherent to DLT generations 1.0 and 2.0 and those that are inherent to DLT 3.0 and beyond. Table 2 describes the range of blockchain risks.

Table 2: **Taxonomy of blockchain risks**

Risks	Source
DLT 1.0 and 2.0	
51% vulnerability	Consensus mechanism
Private key security	Public-key encryption scheme
Criminal activity	Cryptocurrency application
Double spending	Transaction verification mechanism
Transaction privacy leakage	Transaction design flaw
DLT 3.0	
Criminal smart contracts	Smart contract application
Vulnerabilities in smart contracts	Programme design flaw
Under-optimised smart contracts	Programme writing flaw
Under-priced operations	EVM design flaw

Source: (Li, Jiang, Chen, Luo, & Wen, 2017)

The "51% vulnerability risk": The POW-based consensus mechanism embedded within the original Bitcoin blockchain architecture (and other implementations that require POW) are vulnerable to attacks that mobilise more than 50% of the hashing power deployed by miners. Should this happen, attackers can validate inauthentic transactions thus creating validated chains that no longer conform to actual value-

backed transactions. Arbitrarily manipulating and modifying blockchain records would allow attackers to reverse transactions, double-spend the same token, exclude or re-order the history of transactions, interfere with other miners or slow the confirmation of transactions on parallel chains (Li, Jiang, Chen, Luo, & Wen, 2017). All of these are critical risks that would nullify the benefits of blockchain as a way for establishing distributed and decentralised trust.

Part of the 51% risk is a material outcome of the Bitcoin protocol, especially as it relates to the increasing difficulty associated with mining each new coin. As the difficulty of solving cryptographic proof of work has increased, miners have adopted new mining techniques that raise the risk of concentrated hash power.

At the outset, average personal computers were sufficient to mine Bitcoin. However, as the value of Bitcoin tokens increased and the payoff for mining each new token grew, miners started deploying bespoke computing devices optimised and built solely for solving Bitcoin proof of work puzzles. Such application-specific integrated circuit chips (ASICs) have been deployed in large pools to optimise the mining task. They consume significant amounts of power and require investments that favour large institutional miners over smaller ones. The largest mining pools now deploy tens of thousands of ASICs as Bitcoin mining has become industrialised – a risk not anticipated in the Bitcoin protocol.

Figure 14: **Bitcoin hashrate distribution amongst largest mining pools**
(4-day average 20-24 February, 2018)

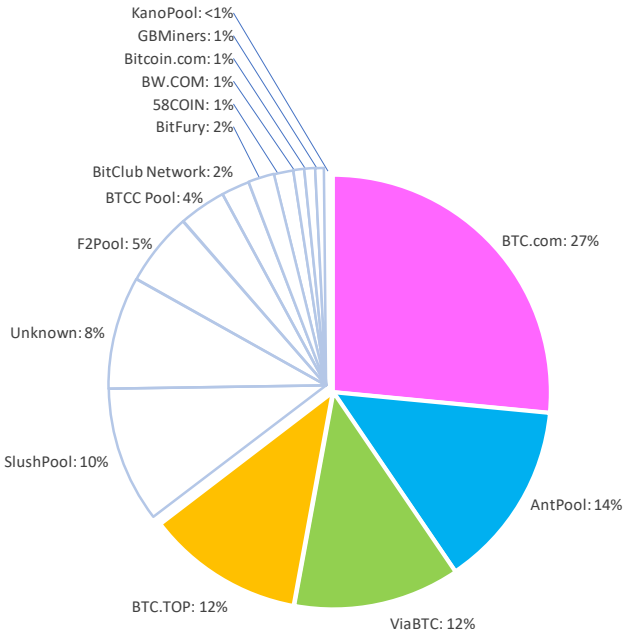

Source: https://blockchain.info/pools

As of early 2018, the top four mining pools account for more than 50% of the overall Bitcoin hashing power (Figure 14). While this concentration does not itself breach the 51% threshold, it raises the potential for collusion amongst large pools to do so or, alternatively, the potential for large miners to be coerced into acting in concert. Both the collisional and coercitative risks are heightened since many large mining pools are spatially concentrated and under the same political regime in China. This concentration stems from the fact that many ASICs are manufactured there (indeed, ASICs manufacturers make up the largest pools) and energy is both plentiful and inexpensive there (Hileman & Rauchs, 2017).

Private key security lies at the heart of most DLT transactions. If the secrecy of the key is compromised, or if it is lost, then core DLT security is compromised since blockchains have no backup in the form of centralised identity management. In the case of compromised keys, or stolen keys, the illegitimate holder of the key can initiate and validate inauthentic transactions.

Though some researchers have discovered vulnerabilities in certain private key signing methods, the cryptographic protection of private keys and their use is generally quite high; compromising or cracking private key encryption remains a daunting task. However, the potential rollout of much more powerful quantum computing will compromise many existing private key encryption algorithms and protocols. There is considerable uncertainty as to when quantum computing will become available but quantum-proof encryption algorithms should already be deployed to anticipate this development. At a minimum, current DLT frameworks should be designed to be upgraded to new cryptographic technologies that will be quantum-proof. At the same time, quantum computing will also provide new tools and security-enhancing functionalities that go far beyond what is possible today.

A more immediate risk is that, just as with physical keys, criminals can and have coerced rightful owners to hand them over under duress. In the case of lost keys, rightful parties can no longer access their transactions, their holdings, or carry out legitimate transactions. Though there are ways to digitally defend against these risks, existing analogue methods to protect against the physical loss or theft of keys will continue to play an important role here (e.g. like keeping a hard copy of the key-password).

Criminal activity: Because of their anonymous, untraceable and decentralised nature, cryptocurrencies built on blockchain have been used in support of criminal activities. This is a risk inherent to cryptocurrencies but not necessarily to the DLTs that underpin them. The fact that cryptocurrencies may be used in criminal transactions should not detract from the potential for DLTs to contribute to better societal outcomes – especially in the case of transport use cases.

Double spending: Consensus mechanisms in blockchain-based cryptocurrencies are designed to limit the risk of double-spending (or double-validation of a single transaction) but this risk is difficult to eliminate. Blockchains that rely on proof-of-work are inherently vulnerable to this type of attack since attackers can use the time it takes to validate a transaction to initiate another transaction using the same token. In the case of blockchain applications in transport, an attacker could initiate one transaction –e.g. accessing a shared vehicle – and then quickly initiate another transaction to access another vehicle under the same identity and using the same access rights. However, unlike double-spending currencies, MaaS-based blockchain transactions will involve physical assets that can have additional security mechanisms built in to prevent fraudulent use.

Transaction privacy leakage: The use of private keys, and especially the use of transaction-specific private keys, helps preserve the anonymity of transactions in blockchain-based cryptocurrency systems. However, the *pattern* of transactions and linking these to other available identifiers can help attackers deduce the actual identity of transaction participants. This is a risk that must be addressed in DLT design, especially as multiple, potentially confounding, data are increasingly available. In this respect, the security of a DLT protocol must extend beyond the specific protocol itself and take into account other systems that it will interact with. This will be especially true in the case of MaaS DLT applications since robust anonymity is already challenging to ensure in present transport use cases (ITF, 2015) (ITF, 2016).

Smart contract vulnerabilities: The ability to embed snippets of executable code in blockchain and other DLT applications is a potential catalyst for widespread uptake of these in MaaS. Nonetheless, the design and use of smart contracts also pose risks that must be addressed in the deployment of the third generation of DLTs for MaaS. The first of these is the potential misuse or malevolent-use of smart contracts themselves – e.g. by designing smart contracts that initiate or support criminal activity. The second risk relates to attacks that exploit poor code and programming errors that exacerbate built-in vulnerabilities in blockchain smart

contracts. Vulnerabilities of these types have already been exploited in known Ethereum smart contract attacks. (Atzei, Bartoletti, & Cimoli, 2017). Poorly optimised smart contracts and under-priced operations can trigger useless and resource-consuming operations that slow or prevent the timely execution of contracts and transaction validation.

Are DLTs safe enough for use in MaaS despite their vulnerabilities? All these vulnerabilities are known and are the focus of much ongoing attention and innovation in the field of cryptography and computer science. To some extent, they are inherent to any set of new computer protocols and the risks they pose are directly related to specific use cases that they support. Clearly, in the case of cryptocurrencies, these vulnerabilities are critical in that they can lead to significant theft and losses.

It is not evident, however, that the risks posed by some of these vulnerabilities are relevant or critical to DLTs used in the case of MaaS. Some of these vulnerabilities could be exploited, for example, to fraudulently access a transport service. The potential loss to operators would, however, be limited and the potential for operators to identify the fraud and initiate corrective and/or punitive action would be high since it would easy to design the system to tightly link actual identity to users. In the end, current transport systems are also vulnerable to fraud, misuse and criminal use. Going forward, the decision to deploy DLTs in support of MaaS should not be predicated on the fact that these present *any* security risks, but rather, that the security risks that they present are *fewer* or *less severe* than those already present in existing systems. Doing so will require a consistent and broad security assessment framework for DLTs that does not yet exist in the field of transport, and MaaS applications in particular.

Data syntax for Mobility as a Service

A common language and data syntax would facilitate "everything-to-everything" connectivity e.g. an internet of mobility, which could underpin DLT-enabled MaaS. This is far from the case today since few of the many transport services available to people use a common data syntax or a shared data referencing framework. This is understandable because such a universal data syntax – an "HTML" of mobility – does not exist. Or at least not in a broad enough form to encompass the wide range of scheduled, un-scheduled, on-demand, peer-to-peer, fixed and free-floating services.

This means that efforts to bring together these services into a common MaaS environment is complicated by the need to convert the various bespoke encoding methods and data formats into interoperable forms. A common language and syntax would allow services to be seamlessly integrated from the planning stages, to the on-trip coordination of the services to the payment and transaction authentication phases. The broad range of data formats also makes it difficult to easily on-board this data into various regulatory and oversight applications on the part of public authorities.

Data incompatibility among service operators

There are historic and institutional reasons for data incompatibility among service providers.

Traditional public transport operators have developed their data solely for internal planning and operations and there has been little need to develop interoperable data formats that can be used by other transport service providers or authorities. Some operators have moved more recently to harmonise at least part of their data on scheduled and real-time service according to a common data syntax for "external" use. Nonetheless, multiple "harmonised" data formats (e.g. VDV 453, VDV 454, SIRI, TRIAS, GTFS, GTFS RT, NetEx) exist and there is no broad consensus on their use that would enable easy, cross-service integration. Cross-service integration, itself, is not necessarily a goal of these operators as some view opening up their schedule or operations data to other for-profit service-delivering parties (especially those producing combined mapping, search and wayfinding services paid for by targeted advertising) as a potential erosion of the value of producing and delivering public transport services.

Mobility data harmonisation and aggregation

Access to public transport data may be negotiated on a case-by-case basis, mediated through an open or permissioned API, and/or provided in direct open access and via an open feed. In many instances, especially in North America, schedule or service sharing is undertaken using the General Transit Feed Specification (GTFS) for scheduled and real-time services. The open nature of GTFS, its simplicity and portability has led it to be quickly adopted by many operators as a default public transport service data encoding format. The broad structure of GTFS has been also used to specify a common referencing system for bike share services - the Global bikeshare feed specification – GBFS. However, neither of these are fully open in that they have closed governance structures as opposed to many DLTs which have open governance structures.

Companies that provide unified information on travel choices and routing options deploy considerable resources collecting, transforming and rendering interoperable multiple types of data. This back-office work is what makes possible inter-modal routing suggestions like those provided by CityMapper, MoovitApple, Baidu and Google Maps.

These companies link into operators' feeds where possible (e.g. public transport, bike and car share and ride services like Didi Chuxing, Uber, Lyft, Grab) and manually complete the rest for relevant markets. This harmonised data is then made accessible in bespoke formats and via proprietary APIs for third-party application developers to use – oftentimes at a cost for large volumes. Having a harmonised and open data structure would potentially erode the control over *which* information is released, (e.g. regarding privileged or paying partners) its *ranking* and *priority* when returned to users, and would potentially lead to revenue losses for some companies under current frameworks. In some cases, like those for transport data integrators and way-finding applications, it would severely challenge current business models.

On the other hand, having a common and harmonised data format on transport services would reduce the information costs associated with coordinating these in broad-scale MaaS applications, which themselves would plausibly generate significant new revenue opportunities and deliver more efficient and optimised transport services. Crucially, it would greatly facilitate the use of data on transport services in a DLT environment.

Data syntax for Mobility as a Service: The Internet of mobility

Such a data syntax should be open and flexible enough to cover the broad range of existing services and to incorporate new services as well – like those that could be delivered by shared, self-driving vehicles. The syntax would ideally be open and harmonised along its broad lines but allow commercial actors to provide proprietary and permissioned-access only data for their commercial service-related data (Figure 15). Such a shared data syntax would be used as a basis for encoding transport services in a blockchain/distributed ledger environment and would thus represent the building blocks for seamless MaaS. This would both enable and underpin the on-the-fly smart contract-clearing transactions that could eventually obviate the need for centralised platforms to deliver MaaS.

Figure 15: **Mobility service data syntax "bins" with open vs. permissioned access layers (indicative)**

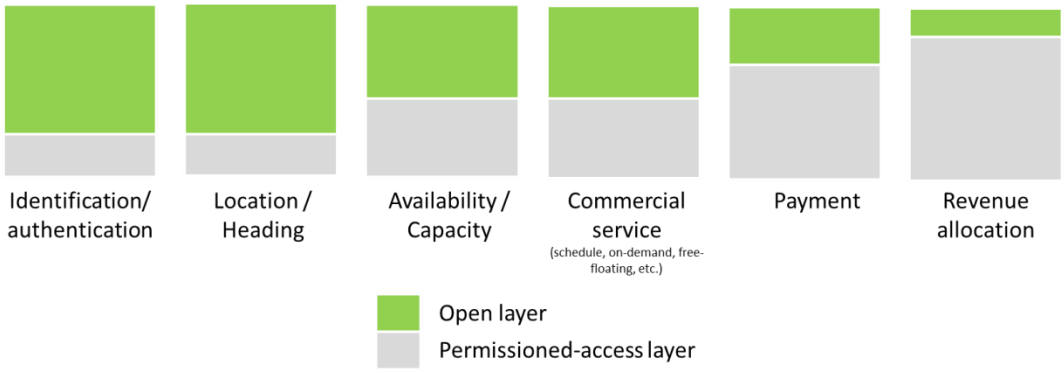

First steps: Minimal open data sharing

Governments can already take action to prepare the ground for increased uptake of DLTs by starting to work on data policies. For instance, new provisions in the Finnish Transport Code (Box 3) already lay the groundwork for a common data structure that could enable seamless MaaS. Rather than focus on data *structure*, the code addresses data *availability* and *usability*. The code calls for transport service providers and regulated entities to establish an open, easily accessible and useable digital channel delivering a common set of data items. These items must include those outlined in Table 3 (Finnish Ministry of Transport, 2017):

Table 3: **Required data reporting elements for passenger transport operators in Finland**

The identity of the service provider, commercial registration number and contact information that a service user can use.
Data regarding the spatial coverage of the service.
Information on payment options.
Information related to the accessibility of the service to those with mobility or other impairments.
Machine-readable information regarding scheduled service operation and spatially-referenced route information.
The location of scheduled traffic stops, stations, terminals with related timetable information.
The period(s) for which the service or timetable information is valid.
For non-scheduled services and for any potential service provider, geospatial information on predetermined stops, stations, terminals, etc.
For non-scheduled or on-demand services, information on the times the services are available.
Information on how to book or hail the services(s) with a link to the booking engine if applicable.
Information on the price of the service including the breakdown into both static and dynamic (e.g. time- or distance-based) fare components, including discounts. This information should allow for cross-service comparison (e.g. for peak hour use).
Dynamic price information and information on available capacity, or a link to the service from which this information is available.
Information regarding restrictions, conditions, extra fees or policies or available options (e.g. regarding baggage transport, policies regarding animals, carriage of children, work stoppages, etc.).
Real time trip planning and en-route data or a link to a service making this information available.
For non-scheduled services, map-based display of the location of available and/or booked vehicles or a link to the service from which the information is available.
Estimates of significant delays or cancellations in services as soon as they are available to service providers.
A link to the web site or other electronic service of the service provider.

These provisions are meant to create an open and level playing field where both small and large operators can more seamlessly coordinate or link their services and create new innovative options or applications. Without being overly prescriptive, these provisions start to set in place a common data framework for MaaS

which could enable more rapid uptake of DLS-based MaaS applications. Given the level of uncertainty around the suitability of first generation DLS for MaaS, requiring minimal and open data sharing seems a prudent and prescient first step.

Common spatial referencing

Describing transport services in a common syntax is not the only use-case where having a common referencing framework is helpful. MaaS services are carried out on streets and have an open framework for creating and sharing street-linked data could also catalyse DLT-enabled MaaS. The SharedStreets referencing system uses topology and other descriptive properties to define locations on streets and transport infrastructure and gets away from some of the constraints inherent in using predefined, proprietary identifiers to describe street locations (SharedStreets, 2018) (ITF, 2018). This referencing system allows multiple users to share street-linked data without sharing underlying proprietary base-map data or even requiring users to agree on a common base map. This form of open spatial syntax can also catalyse MaaS applications, especially when bundled into DLT-based applications.

Box 3: **Finnish Act on Transport Services**

The Finnish Transport Code reform centres on revising the Act on Transport Services. This Act brings together legislation on transport markets and creates preconditions for digitalisation and new business models in transport. Its core aim is provision of customer-oriented transport services. The provisions relating to the introduction of Intelligent Transport Systems linked to the ITS Directive entered into force on 1 October, 2017 and those relating to the interoperability of data and information systems , on 1 January, 2018. The other provisions will enter into force on 1 July, 2018.

The new Act creates a framework for a more efficient arrangement of publicly subsidised passenger transport by using digitalisation, combined transport and different fleet types. The objective set in the Government Programme is to achieve a 10% saving in publicly subsidised passenger transport from 2017. Implementation of the objectives of the Act on Transport Services requires the opening of data and the handling of matters through open interfaces. Opening of data is continued with regard to data on the use of mobility services. The Finnish Transport Agency would be obligated to open data received on the use of services through open interface, in a form where it cannot be linked to individual users, service providers or services.

The Act also lays down provisions for the interoperability of ticket and payment systems. Offering trip chains and combined services would be eased by enabling acting on another's behalf, for instance, the provider of a combined service could incorporate tickets for all modes of transport, car hire service, various serial and seasonal products as well as discounts of a combined mobility service by acting on the customer's wishes or on the customer's behalf in different services.

A third tranche of Transport Code reform is currently underway addressing the regulation of logistics and freight services.

Figure 16: **Transport Code regulation of logistics and freight services in Finland**

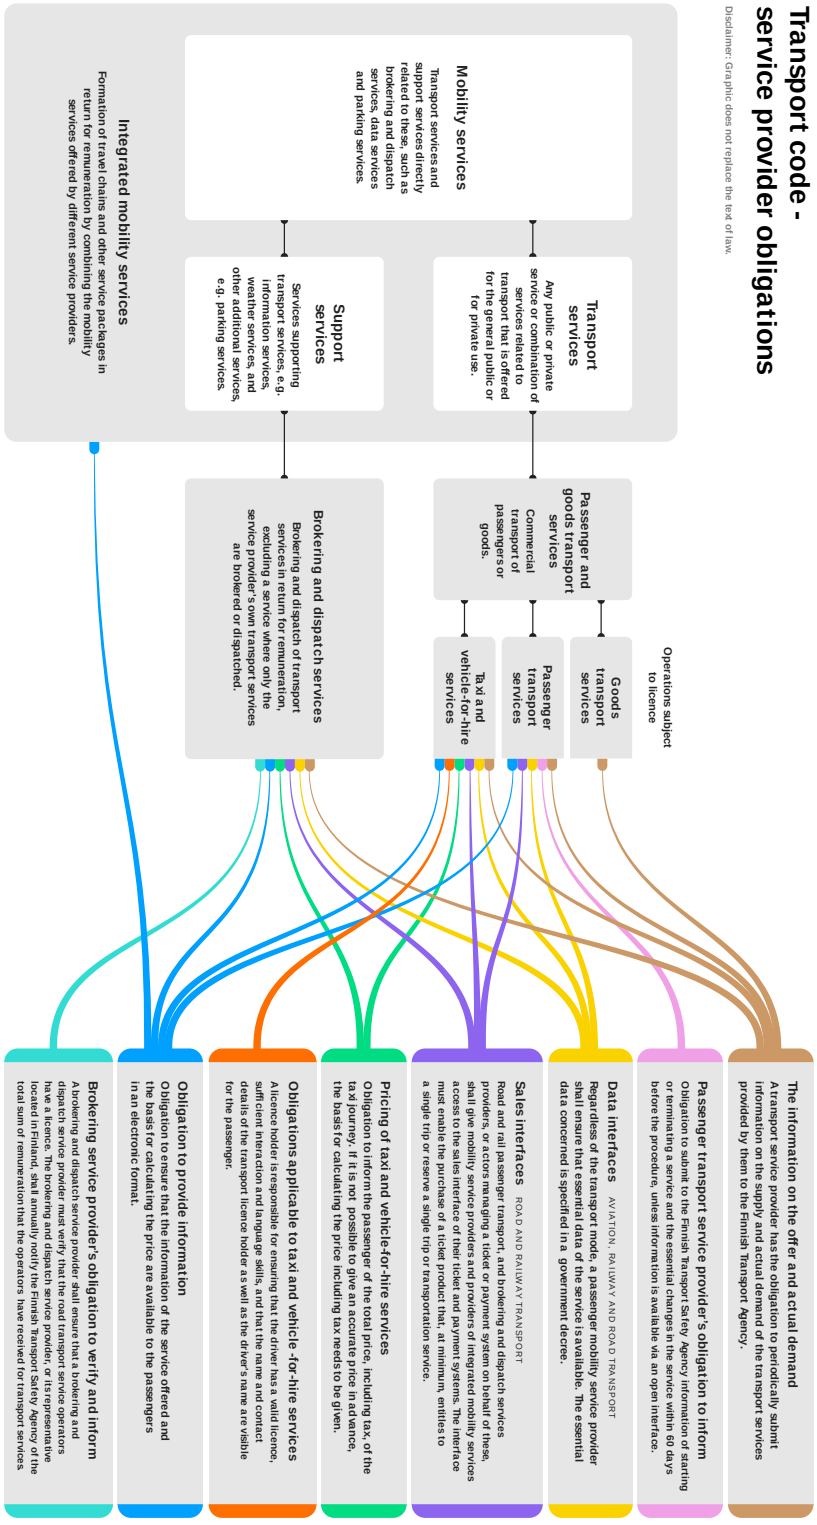

Source: Finnish Ministry of Transport and Communications.

Open algorithms and other alternatives to data sharing

It is somewhat ironic that all the while tremendous amounts of new, granular and potentially useful data are generated by transport, *access* to this data among the actors who regulate urban transport is increasingly complicated. There are technical reasons for this inherent in the range of data formats involved as described in the previous section, but, for the most part, the barriers to sharing this data lies more broadly in the commercial sensitivity of the information contained and in the privacy implications for travellers whose activity often generates the data.

Data sharing and personal data

Against this backdrop of a shift in data accessibility is a broader, more fundamental question that relates to the very nature of personal data ownership and the extent with which individuals own or can otherwise control the use of data that pertains to them. This is an area where there is a broad split in practice and in law. Within the European Union, the General Data Protection Regulation (GDPR), which enters into force in May 2018 and becomes fully enforceable two years thereafter, sets out clear answers to those questions (Box 4). In the case of personal data, the law states that:

"Personal data is any information that relates to an identified or identifiable living individual. Different pieces of information, which collected together can lead to the identification of a particular person, also constitute personal data." (European Commission, 2018)

Crucially, in the case of transport-specific applications, this definition also includes geo-spatial and locational data. This definition includes data that has been de-identified, encrypted or pseudonymised but could be re-identified directly or using confounding data (ITF, 2015). Data that cannot be re-identified or de-anonymised falls outside of the scope of the law.

The GDPR is clear on the right of individuals to control the collection, use and dissemination of their personal data. It also underscores the right for individuals to discover what data is held on them and to be able to easily and conveniently transfer this data amongst service providers.

Data protection rules are different in many other jurisdictions and are generally less extensive or weaker than those in the GDPR though there are exceptions – South Korea's Personal Information Protection Act of 2011 is well aligned with the scope and strength of the GDPR and Japan's Personal Information Protection Act of 2015 converges with many aspects of the EU GDPR.

The emerging legal frameworks around the protection of personal data open the door to new models for handling and processing personal data, including the data that will be created and, in turn, will fuel seamless MaaS applications. Blockchain and DLTs are uniquely suited to personal privacy-preserving applications. Approaches such as the MyData model developed by a consortium of private and public actors in Finland have promise as a way forward.

Box 4: **The EU General Data Protection Regulation**

The EU General Data Protection Regulation (GDPR) seeks to protect all EU citizens from privacy and data breaches in an increasingly data-driven world that is vastly different to when the prior 1995 directive was established. Although the key principles of data privacy still hold true to the previous directive, many changes have been proposed to the regulatory policies; the key points of the GDPR are below.

Increased Territorial Scope (extra-territorial applicability): Arguably the biggest change to the regulatory landscape of data privacy comes with the extended jurisdiction of the GDPR, as it applies to all companies processing the personal data of data subjects residing in the Union, regardless of the company's location. The GDPR will apply to the processing of personal data by controllers and processors in the EU, regardless of whether the processing takes place in the EU or not. The GDPR will also apply to the processing of personal data of data subjects in the EU by a controller or processor not established in the EU, where the activities relate to: offering goods or services to EU citizens (irrespective of whether payment is required) and the monitoring of behaviour that takes place within the EU.

Penalties: Under GDPR organisations in breach of GDPR can be fined up to 4% of annual global turnover or EUR 20 million (whichever is greater).

Consent: The conditions for consent have been strengthened, and companies will no longer be able to use long illegible terms and conditions full of legalese, as the request for consent must be given in an intelligible and easily accessible form, with the purpose for data processing attached to that consent. Equally, consent must be clear and distinguishable from other matters and provided in an intelligible and easily accessible form, using clear and plain language. It must be as easy to withdraw consent as it is to give it.

Data Subject Rights:

Breach Notification: breach notification will become mandatory in all member states where a data breach is likely to "result in a risk for the rights and freedoms of individuals". This must be done within 72 hours of first having become aware of the breach. Data processors will also be required to notify their customers, the controllers, "without undue delay" after first becoming aware of a data breach.

Right to Access: Part of the expanded rights of data subjects outlined by the GDPR is the right for data subjects to obtain from the data controller confirmation as to whether or not personal data concerning them is being processed, where and for what purpose. Further, the controller shall provide a copy of the personal data, free of charge, in an electronic format.

Right to be Forgotten: The right to be forgotten entitles the data subject to have the data controller erase his/her personal data, cease further dissemination of the data, and potentially have third parties halt processing of the data.

Data Portability: GDPR introduces data portability - the right for a data subject to receive the personal data concerning them, which they have previously provided in a 'commonly used and machine readable format' and have the right to transmit that data to another controller.

Privacy by Design: Privacy by design as a concept has existed for years now, but it is only just becoming part of a legal requirement with the GDPR. At its core, privacy by design calls for the inclusion of data protection from the onset of the designing of systems, rather than an addition.

Adapted from: https://www.eugdpr.org/the-regulation.html

In the MyData model, individuals retain rights to their data (in compliance with the GDPR) and exercise control over its use. They can share the data or allow its transfer to the digital services or platforms they designate and trust. To do so, they leverage the parametric privacy built into the protocol that helps allocate and track the use of personal data (Finnish Ministry of Transport and Communications, 2017).

Box 5: **MyData**

The MyData Alliance is an open community, which advance MyData pilots and share knowledge and resources. The aim is to develop national, internationally scalable, interoperability models for personal data management. At the heart of the MyData approach are three foundational principles:

Human centric control and privacy: Individuals are empowered actors, not passive targets, in the management of their personal lives both online and offline – they have the right and practical means to manage their data and privacy.

Usable data: It is essential that personal data is technically easy to access and use – that it is accessible in machine readable open formats via secure, standardised APIs (Application Programming Interfaces). MyData is a way to convert data from closed silos into an important, reusable resource. It can be used to create new services which help individuals to manage their lives. The providers of these services can create new business models and economic growth for society.

Open business environment: Shared MyData infrastructure enables decentralised management of personal data, improves interoperability, makes it easier for companies to comply with tightening data protection regulations, and allows individuals to change service providers without proprietary data lock-ins.

MyData is an infrastructure-level approach for ensuring data interoperability and portability. It is sector independent – there is currently significant progress being made in individual sectors, such as health and finance, but a cooperative approach across all sectors has more promise. Finally, it allows consent-based data management and control.

The aim is to provide individuals with the practical means to access, obtain, and use datasets containing their personal information, such as: purchasing data, traffic data, telecommunications data, medical records, financial information and data derived from various online services. Organisations holding personal data are encouraged to give individuals control over this data, extending beyond their minimum legal requirements to do so.

The MyData architecture is based on interoperable and standardised MyData accounts. The account model provides individuals with an easy way to control their personal data from one place even while the data is created, stored, and processed by hundreds of different services. For developers, the account model facilitates access to data and removes dependencies on specific data aggregators.

Interoperability within the data management system can be understood as functioning similarly to interoperability in mobile telephone networks. Both systems require a common network that connects distributed nodes. Global interoperability and transferability of MyData accounts (and thus individual's consents) between operators requires further standardisation and design e.g. on trust networks, data formats, and semantics.

As such, the MyData model is complementary to the EU General Data Protection Rules (GDPR). It also has the potential to integrate several DLT applications in its component parts – e.g. for "consent-tagging" data within a permissioned DLT – and ensuring trust and high levels of cryptographic security for consent-based access to personal data.

Source: Adapted (Finnish Ministry of Transport and Communications, 2018)

New models for gaining trusted insight from privately held data

Transport-related data collection, knowledge and insight are increasingly shifting away from the public sector and into the private sector. This shift implies a growing information asymmetry between those in charge of regulating transport activity and public space and those with actionable and relevant information to do so. This has led to public authorities either purchasing data from commercial actors or compelling them to provide their data. Neither approach ultimately satisfies both parties and yet "*sell me your data*" or "*give me your data*" largely comprise the only two data discovery options considered by the public sector.

One alternative to the current data "ask" is to entrust and house data with neutral third-parties that mediate access to the data or its analytic outputs according to rules agreed by all. This third party could be a university or a dedicated public agency (though public universities and agencies may face similar conflicts

as transport authorities in legal regimes where public authorities must adhere to "right-to-know" laws) or an audited commercial data-holding operator. This model is starting to be explored with some agencies.

Recent changes in data science and new alternatives to data sharing provide new ways of extracting useable insight from raw data. These new methods could form the basis of a new data deal amongst commercial services, individuals and public authorities. These approaches essentially dematerialise the trusted third party model into trusted and vetted code operations that are transparent for all parties to understand, see and authorise.

In traditional data-sharing approaches, data itself is transmitted from where it is collected and housed to a commercial partner or to a public agency – with all of the competition and privacy risks that this might entail. That is because having the data in hand has been the best way to ensure the correctness, veracity and trustworthiness of the analytical outputs based on the data. However, rather than relying on transmitting *data* between parties, new emerging approaches rely on trading trusted and vetted *code* – essentially transmitting code to the original data source and executing its analysis there and allowing these algorithms to run analytic operations on, and return trusted responses from, remotely-held data.

Figure 17: **Advantages of transmitting code instead of data (or vice-versa)**

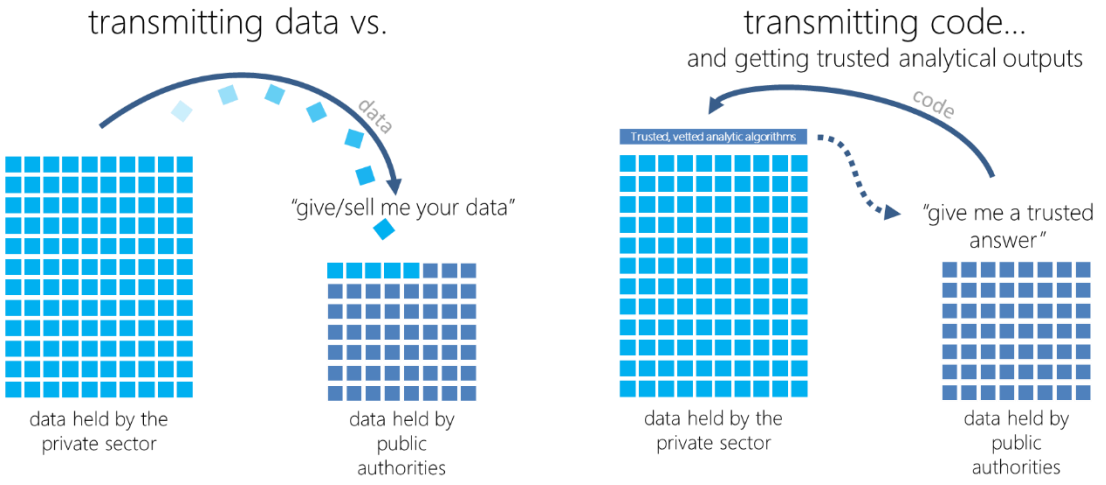

OpenTraffic, a data-analytical framework built on the concept of code-sharing instead of data-sharing has been developed by the World Bank and partners in the context of the Open Transport Partnership. OpenTraffic has been implemented in partnership with Grab, a major ride-service platform operator in Southeast Asia. With OpenTraffic, authorities receive trusted information regarding traffic speeds derived from algorithms working directly on Grab's servers behind the company's firewall. Raw data is never transmitted outside of the company but because both Grab and public authorities have vetted the algorithms operating on the company's servers, the specific output concerning traffic speeds is considered to be correct and trustworthy (Sharpin, Adriazola-Steil, & Canales, 2017).

The OpenTraffic approach is also at the heart of several other initiatives seeking to find alternative methods of extracting trustworthy insight from privately- or commercially-held data beyond traditional data-sharing. One of the first of these is the "Safe Answer" framework developed at MIT (de Montjoye, Shmueli, Wang, & Pentland, 2014). This mechanism is built around data users submitting code snippets that mediate on individuals' raw data in their personal "data store" without releasing any of that data itself. Under a "personal data store" framework, the Safe Answers approach calls for potential data users to submit a request for information regarding an individual's data. The question could be "is this person close to my store?", "how much time does this person spend in traffic on a weekday?" or "does this person use the

underground on weekends? If the individual accepts that request (perhaps granting this acceptance in return for a service or other form of compensation from the data user), the data user submits a standardised snippet of code that then interacts with that person's personal data store, querying GPS log data, accelerometer data, or other form of location/trajectory data required to answer the question. The answer is sent back to the data user without sensitive location or trajectory data ever having been divulged.

Another approach extending the "Safe Answer" framework to companies is embodied in the Open Algorithm (OPAL) project at MIT (OPAL, 2017). OPAL's core consists of an open platform and algorithms that run on the servers of partner companies behind their firewalls to extract key indicators of relevance for a wide range of potential users. This approach could, for example, return from ride-service, taxi and public transport operators an aggregate density of pick-up and drop-off events at a block face level to authorities concerned about traffic congestion and safety without ever revealing sensitive data from operators.

Yet another approach that builds on the principles of bringing code to the data rather than data to the code is the Enigma framework, also developed at MIT. Enigma is a decentralised computation platform with guaranteed privacy. The privacy of Enigma is ensured by data queries being computed in a distributed way, without a trusted third party. Data is split between different nodes, and they compute functions together without leaking information to other nodes. Specifically, no single party ever has access to data in its entirety; instead, every party has a meaningless (i.e., seemingly random) piece of it.

Enigma is also scalable. Unlike blockchains, computations and data storage are not replicated by every node in the network. Only a small subset per x forms each computation over different parts of the data. The decreased redundancy in storage and computations enables more demanding computations. The significant new feature Enigma introduces is the ability to run computations on data, without having access to the raw data itself. Enigma provides an alternative to data-sharing frameworks in that traditional data-sharing is an irreversible process; once it is sent, there is no way to take it back or limit how it is used. Allowing access to data for secure computations is reversible and controllable, since no one but the original data owner(s) ever see the raw data. This presents a fundamental change in current approaches to data analysis.

OpenTraffic, Safe Answers, OPAL and Enigma all provide examples of novel ways of applying data science to overcome some of the challenges inherent in traditional data-sharing frameworks. These models are, like Blockchain and DLT, at their very early stages of development (for OPAL and Enigma) and implementation (for Safe Answers and OpenTraffic) and yet they outline a future direction for actionable insight to be derived from privately held and commercially sensitive data. The broader implementation of these or similar approaches can help ensure lighter regulatory frameworks while ensuring better regulatory outcomes and can facilitate broad-scale cooperation amongst multiple, sometimes competing, partners in the context of scaled-up MaaS applications.

What policies for now, what principles for later?

Distributed ledgers as both a technology and as a foundation for new business and regulatory processes are still very much in their early days and are evolving in ways that are hard to predict. This complicates the task for public authorities to critically assess what role they will play in Mobility as a Service (MaaS) and what role, if any, there will be for public authorities in transport-relevant DLT deployment. Clearly much that is yet unanticipated will play a role on the pathways for DLTs to be an enabler of more seamless and interconnected transport services. How then should public authorities establish policies today and develop longer-term principles to guide action in the future? A starting point is to think about the deployment of DLTs in transport as a continuum, not as an endpoint. Figure 18 illustrates the near and mid-term perspectives for blockchain applications in transport as proposed by a working group convened by the French Ministry of Transport in 2017.

Figure 18: **Mobility-related use cases for distributed ledgers**

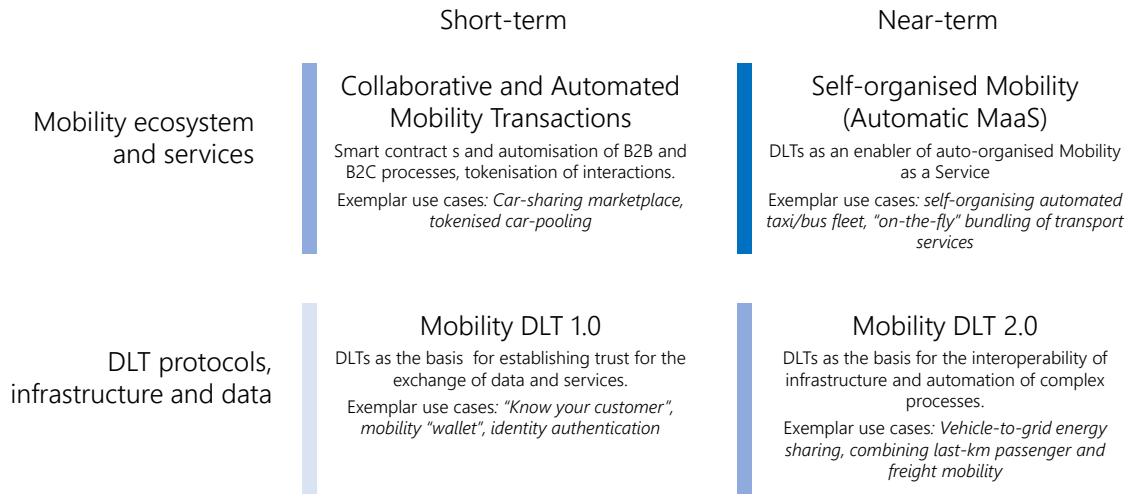

Source: Adapted from Ministère de la transition écologique et solidaire, 2017

Authorities should already be readying themselves and starting to support the use of DLTs for establishing trust and identity in their own systems and applications – generally testing Mobility DLT 1.0 applications. The next steps should look to ways in which complex processes can be automated in a trustworthy way within open ecosystems. As authorities gain confidence and improve their regulatory skills, and as both they and the private sector deploy more complex and interlinked use cases, DLTs can increasingly form the basis for a largely self-organised open mobility ecosystem – automatic MaaS. This transition will imply a change from business-as-usual for both the public and private sector. The following points should help the first steps of this transition.

Focus on DLTs, not on cryptocurrencies: Uncertainty regarding DLTs is set against a backdrop of increasing chatter about blockchains and their use as the foundational technology in various cryptocurrencies (Figure x). The fact that much of the focus on blockchain has been in the field of cryptocurrencies comes with a risk that public authorities may evaluate and judge the suitability of the technology solely on the basis of its deployment in (potentially risky) applications. This would be unfortunate since, as noted by Derek Thompson, senior editor at The Atlantic, "[Bitcoin is] for now, a frankly terrible currency built on top of a potential transformative technology" (Thompson, 2017). The point isn't if current DLT applications are suited to applications like MaaS but if the foundational principles are

robust enough for the MaaS (or any other transport sector-related) task. If so, and it seems likely they are, then technology will surely evolve and therefore policy should anticipate this development.

At the same time, despite multiple potential use cases and business models, there are very few fully scaled-up examples of non-currency blockchain implementations in any domain with the exception of the Estonian KSI e-identity blockchain.

Build the building blocks for DLT-enabled MaaS: Rather than supporting broad-scale deployment of existing DLTs, public authorities should ensure that the necessary building blocks are in place for future DLTs. These could include harmonised identifiers, a shared and common data syntax in support of the internet of mobility and a regulatory framework that anticipates future DLT developments. These are generally applications in the "Mobility DLT 1.0" quadrant of Figure 18.

Task-based rather than MaaS-based DLTs for now: As noted in this report, the use of distributed ledger technologies to support a seamless, peer-to-peer distributed and open mobility-as-a-service ecosystem could be one application of DLTs. Nonetheless, the complexity of such an ecosystem, and the need to ensure buy-in from a wide number of market and public actors might act against early adoption of the technology for this particular use case despite many early proof-of-concept initiatives. Building up to the "Self-organised Mobility" quadrant of Figure 18 will first require all actors, public and private, to develop experience and proficiency in specific DLT-supported tasks before they can be assembled into a single ecosystem.

DLTs are ready for "slow and (relatively) small" but not yet for "big and fast": Current blockchain implementations are limited because they fail to scale and are relatively slow. Nonetheless, they are still suited for MaaS tasks that are not sensitive to limitations in volume or speed. These include identity management, licensing and registration and asset tracking. These use cases can serve as a test bed that will allow stakeholders to become familiar with DLT-supported MaaS applications. MaaS tasks that require processing more real-time logging and high-volume data will require newer DLT models purpose-built for speed and "internet of things" applications. This is especially true for anything involving recording of current state that is essential for further MaaS tasks – like dispatching. Certain technologies, like DLTs based on directed acyclical graphs, can already be tested although their large-scale uptake for MaaS may not be immediate.

Open but permissioned DLT deployment aligned with other regulatory roles: A fully open and permissionless DLT model like the Bitcoin blockchain protocol seems ill-suited for many MaaS tasks – especially those that require low latency or must process in real-time. At the outset, a limited, open but permissioned DLT model may be better suited for early MaaS applications. In this type of implementation, all potential service providers could join but credentialing and onboarding would be managed with centrally controlled permissions – much along the lines of licensing and registration requirements currently in existence today. The governance of the MaaS DLT protocol and validation methods may also fall under a central authority.

This raises the question of who manages that ledger and what are the governance rules that establish consensus protocols and validation authority? This seems to be an area where public authorities may wish to focus early and anticipatory policy action – either to establish a MaaS DLT ecosystem under a common and public authority-led governance framework, or, to establish a set of minimum core governance rules that DLTs in the MaaS space must adhere to.

Both approaches beg the question "what is the rationale for public policy intervention in MaaS-related DLTs?" This is an open question, but governments typically intervene in domains where the cumulative actions of individuals and businesses, each optimising their own interests, lead to sub-optimal outcomes from a societal perspective.

Machine-readable laws and regulations: Much of the regulation of transport stems from a desire to minimise externalities and market failures that counter publicly desired outcomes. These include anti-competitive behaviour, excessive congestion, environmental damage, poor safety and inequitable access. In an era where more and more transport activity will be the combined result of peoples' decisions and the way in which this agency is guided, facilitated or generated by the underlying web of code-based algorithms and protocols that mediate a growing share of transport activity.

At a minimum, public policy should understand the algorithmic substrate of transport. But this is likely not sufficient, and policy may be needed to guide transport and societal outcomes towards those outcomes that public authorities have a mandate to deliver. And this is precisely where there is a real and growing disconnect between how governments have traditionally regulated activities and how these activities actually operate.

Much of the classic regulatory framework is built around a set of rules embodied in the legislative framework. The code of law and regulation is analogue, paper-based and crafted in human language. This has been the right framework in a world where most of the decisions giving rise to an activity have been made by people and corporations. But in a world increasingly characterised by the outcomes of algorithms embodied in code and software, this may no longer be sufficient. In the case of DLTs, ensuring that the technical code that defines the rules and governance principles for distributed ledgers has the same rigour as legislative code is important (UK Government Office for Science, 2016).

The successful deployment of DLTs in support of MaaS will necessitate a governance framework that ensures value creation for participants while protecting their interests, those of society at large and defending the system from systemic risks or criminal activity. Here, there are emerging opportunities to build on interactions between legal and technical code. Rather than relying solely on "analogue" code to guide public outcomes, technical code could be integrated into the regulatory process to ensure the delivery of public regulatory outcomes.

To be clear, understanding and ensuring that technical code and legislative code work together to deliver desired societal outcomes does not mean that authorities should vet algorithms and approve or disapprove technical solutions developed by the private sector. This would likely be counter-productive and innovation-reducing. Rather, governments should start to think about the ways in which they can translate legislative code and the principles it embodies into a format and a framework that can be easily integrated into algorithmic decision-making. Governments should ensure that they can directly monitor the outcome of different technical systems on public policy objectives (Box 6). This will require a fundamentally new approach and regulatory skillset than has traditionally been deployed by transport authorities. Other government sectors have already started to operate this transition towards "RegTech", most notable financial oversight authorities, but much has yet to be invented in order to build a robust technical-legislative framework for digital actors in transport.

Box 6: **Algorithmic Impact Assessment**

As automated decision-making systems increasingly work their way into the public and private spheres, including in transport (route-finding, automated driving, ride-service dispatching, etc.), authorities have come under increasing pressure to assess the impact of these systems on people and public policy goals. At a minimum, public authorities should include this assessment for the decisions they make and the decision support tools they employ. These algorithmic impact assessments should:

- Respect the public's right to know which systems impact their lives and how they do so by publicly listing and describing algorithmic systems used to make significant decisions affecting identifiable individuals or groups, including their purpose, reach, and potential public impact.
- Ensure greater accountability of algorithmic systems by providing a meaningful and ongoing opportunity for external researchers to review, audit, and assess these systems using methods that allow them to identify and detect problems.
- Increase public agencies' internal expertise and capacity to evaluate the systems they procure, so that they can anticipate issues that might raise concerns, such as disparate impacts or due process violations; and;
- Ensure that the public has a meaningful opportunity to respond to and, if necessary, dispute an agency's approach to algorithmic accountability. Instilling public trust in government agencies is crucial—if the algorithmic impact assessment (AIA) doesn't adequately address public concerns, then the agency must be challenged to do better.

Beyond its use for government decision-making, there may be scope to assess whether AIA methodologies could also be adapted for use more broadly in the regulation of transport services and technologies.

Source: Adapted from: (AI Now Institute, 2018)

Bibliography

(n.d.). Retrieved from https://en.wikipedia.org/wiki/ARPANET

(n.d.). Retrieved from https://en.wikipedia.org/wiki/Fax

AI Now Institute. (2018, February 21). *Algorithmic Impact Assessments: Toward Accountable Automation in Public Agencies*. Retrieved March 18, 2018, from Medium: https://medium.com/@AINowInstitute/algorithmic-impact-assessments-toward-accountable-automation-in-public-agencies-bd9856e6fdde

Amoros, R. (2018, January 10). *Transactions Speeds: How Do Cryptocurrencies Stack Up To Visa or PayPal?* Retrieved January 12, 2018, from howmuch.net: https://howmuch.net/sources/crypto-transaction-speeds-compared

Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A Survey of Attacks on Ethereum Smart Contracts SoK. *Proceedings of the 6th International Conference on Principles of Security and Trust - Volume 10204.* New York: Springer-Verlag.

Blockchain.com. (2018). *Average Confirmation Time*. Retrieved February 20, 2018, from https://blockchain.info/charts/

Blockchainfirst. (2017). *The first Multipurpose Blockchain enabled EV Charging Station*. Retrieved January 23, 2018, from Medium: https://medium.com/@blockchainfirst/the-first-multipurpose-blockchain-enabled-ev-charging-station-d8265c1bcb38

Casey, T., & Valovirta, V. (2016, March). Towards an open ecosystem model for smart mobility services: The case of Finland. Retrieved February 11, 2018, from http://www.vtt.fi/inf/pdf/technology/2016/T255.pdf

Croman, K., Decker, C., Eyal, I., Gencer, A., Juels, A., Kosba, A., . . . Wattenhofer, R. (2016, February 26). On scaling decentralised Blockchains: A position paper. *Financial Cryptography and Data Security: FC 2016 International Workshops*. Christ Church, Barbados.

Dardayrol, J.-P. (2017, August). Blockchains et smart contracts : des technologies de la confiance ?

de Montjoye, Y.-A., Shmueli, E., Wang, S., & Pentland, A. (2014, July 9). openPDS: Protecting the Privacy of Metadata through SafeAnswers. *PLOS one*. Retrieved May 4, 2018, from http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0098790

DOVU. (2018). *DOVU: Blockchain-powered mobility*. Retrieved January 10, 2018, from dovu.io: https://dovu.io

e-estonia. (2018). *we have built a digital*. Retrieved April 27, 2018, from e-estonia.com: https://e-estonia.com/

Estonian Tax and Customs. (2017, June 1). *Income declared through Uber and Taxify has overwhelmingly increased*. Retrieved December 12, 2017, from Republic of Estonia Tax and Customs Board: https://www.emta.ee/eng/income-declared-through-uber-and-taxify-has-overwhelmingly-increased

Ethereum. (2018). *Sharding FAQ*. Retrieved February 18, 2018, from https://github.com/ethereum/wiki/wiki/Sharding-FAQ

Etherisc. (2017, November 23). *Blockchain to Automate Validation of Car Insurance Policies in Mexico*. Retrieved January 12, 2018, from Medium.com: https://blog.etherisc.com/https-medium-com-etherisc-blockchain-to-automate-validation-of-car-insurance-policies-in-mexico-ea7a72e87bf2

Etherscan. (2018). *Ethereum Average Blocktime Chart*. Retrieved February 14, 2018, from https://etherscan.io/chart/blocktime

European Commission. (2018). Retrieved from https://ec.europa.eu/info/law/law-topic/data-protection/reform/what-personal-data_en

EY. (2017, November 16). *EY expands European teams' presence to advance blockchain development worldwide*. Retrieved January 10, 2018, from www.ey.com:

http://www.ey.com/gl/en/newsroom/news-releases/news-ey-expands-european-teams-presence-to-advance-blockchain-development-worldwide

Finnish Ministry of Transport and Communications. (2017). *Analysis of the application areas and applications of blockchain technolofgy in transport.* Helsinki: LVM.

Finnish Ministry of Transport and Communications. (2018). *A Nordic Model for human-centered personal data management and processing.* Helsinki: Finnish Ministry of Transport and Communications. Retrieved May 2, 2018, from https://www.lvm.fi/documents/20181/859937/MyData-nordic-model/2e9b4eb0-68d7-463b-9460-821493449a63?version=1.0

Hackett, R. (2018, January 16). *IBM and Maersk Are Creating a New Blockchain Company*. Retrieved February 12, 2018, from fortune.com: http://fortune.com/2018/01/16/ibm-blockchain-maersk-company/

Hileman, G., & Rauchs, M. (2017). *Global Cryptocurrency Benchmarking Study.* Cambridge: Cambridge Centre for Alternative Finance.

IOTA. (2018). *IOTA*. Retrieved March 7, 2018, from iota.org: https://iota.org

IOTA is centralized. (n.d.). Retrieved from https://medium.com/@ercwl/iota-is-centralized-6289246e7b4d

ITF. (2015). *Big Data and Transport: Understanding and Assessing Options.* Paris: International Transport Forum at the OECD.

ITF. (2016). *Data-Driven Transport Policy.* Paris: International Transport Forum at the OECD.

ITF. (2018). *The Shared-Use City: Managing the Curb.* Paris: International Transport Forum at the OECD.

ITIC. (2018, January 4). *IOTA Partners with ITIC to Build A Global Alliance of Smart Mobility Testbeds*. Retrieved March 16, 2018, from ITIC: http://www.itic-sc.com/iota-partners-with-itic-to-build-a-global-alliance-of-smart-mobility-testbeds/

Kapsch TraffiCom. (2017). *Quontoz Blockchain Tech: Payment and Transaction processes proof-of-concept*. Retrieved January 23, 2018, from Factory: Accelerating with Kapsch TraffiCom: http://www.factory1.net/

Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2017). A survey on the security of blockchain systems. *Future Generation Computer Systems*. doi:https://doi.org/10.1016/j.future.2017.08.020

Lightening Network. (2018). *Lightening Network: How it works*. Retrieved February 15, 2018, from lightening.network/how-it-works/

Ministère de la transition écologique et solidaire. (2017, December). Rapport du groupe de travail blockchain. *Synthèse des ateliers de l'innovation*. (C. Kremer, Ed.) Retrieved March 5, 2018, from https://www.assisesdelamobilite.gouv.fr/file/1472/download?token=5gms1aI-

MOBI. (2018). *mobility open blockchain initiative*. Retrieved May 4, 2018, from www.dlt.mobi

Narula, N. (2017, SEptember 7). *Cryptographic vulnerabilities in IOTA*. Retrieved November 22, 2017, from Medium.com: https://medium.com/@neha/cryptographic-vulnerabilities-in-iota-9a6a9ddc4367

NYU. (2018). *What is Data Science?* Retrieved February 2, 2018, from NYU University-wide initiative in data science: https://datascience.nyu.edu/what-is-data-science/

OMOS. (2017). *Open Mobility System*. Retrieved January 10, 2018, from OMOS: www.omos.io

OPAL. (2017). *Open Algorithms (OPAL) Project*. Retrieved May 4, 2018, from OPAL: http://www.opalproject.org/about-us/

Parkgene. (2017). *Parkgene*. Retrieved May 3, 2018, from parkgene.io/: https://parkgene.io/

Pavel, I. (2017, August). La Blockchain -- Les defis de son implémentation. *Réalités Industrielles*.

Polis. (2017, September 4). Mobility as a Service: Implications for Urban and Regional Transportation. *Discussion Paper*. Bruxelles: Polis.

Porsche. (2018, 2 22). *Porsche introduces blockchains to cars*. Retrieved March 4, 2018, from Porsche: https://newsroom.porsche.com/en/themes/porsche-digital/porsche-blockchain-panamera-xain-technology-app-bitcoin-ethereum-data-smart-contracts-porsche-innovation-contest-14906.html

Raiden Network. (2017). *Raiden Overview*. Retrieved February 20, 2018, from https://raiden.network/faq.html

Share & Charge. (2017). *Share and Charge*. Retrieved April 20, 2018, from shareandcharge.com: http://shareandcharge.com/

SharedStreets. (2018). *A shared language for the world's streets*. Retrieved March 28, 2018, from SharedStreets: www.sharedstreets.io

Sharpin, A., Adriazola-Steil, C., & Canales, D. (2017, July 17). *'Open Traffic' Provides Unprecedented Data to Urban Policymakers*. Retrieved January 17, 2018, from The CityFix: http://thecityfix.com/blog/open-traffic-provides-unprecedented-data-to-urban-policymakers-anna-bray-sharpin-claudia-adriazola-steil-diego-canales/

Stöcker, C. (2017, June 24). *Implementing first Industry 4.0 Use Cases with DAG Tangle—Machine Tagging for Digital Twins*. Retrieved April 3, 2018, from Medium.com: http://bit.ly/2isSqAQ

Summermann, D., Oge, C., Smolenski, M., Fridgen, G., & Rieger, A. (2017, December). Open Mobility System (OMOS) Concept Paper. MotionWerk GmbH, Fraunhofer FIT.

swytch. (2017). *Swytch: helping th eworld reduce its carbon footprint*. Retrieved May 2, 2018, from swytch.io: https://swytch.io/

The Transparency Compendium. (n.d.). Retrieved from https://blog.iota.org/the-transparency-compendium-26aa5bb8e260

Thompson, D. (2017, November 30). *Bitcoin Is a Delusion That Could Conquer the World*. (T. Atlantic, Editor) Retrieved February 15, 2018, from www.theatlantic.com: https://www.theatlantic.com/business/archive/2017/11/bitcoin-delusion-conquer-world/547187/

Transport Systems Catapult. (2016, July). Mobility as a Service: Exploring the opportunity for Mobility as a Service in the UK. United Kingdom.

TTio Protocol. (2017). *TSio Protocol: unifying transport with interoperable mobility infrastructure*. Retrieved January 10, 2018, from tsioprotocol.com: https://tsioprotocol.com/whitepaper/

UK Government Office for Science. (2016). Distributed Ledger Technology: Beyond Blockchain. (M. Peplow, Ed.)

Vinchain. (2017). *Vinchain*. Retrieved February 12, 2018, from vinchain.io/: https://vinchain.io

ZF. (2017). *ZF, UBS and innogy Innovation Hub Announce the Jointly Developed Blockchain Car eWallet*. Retrieved January 12, 2018, from ZF Press Center: https://press.zf.com/site/press/en_de/microsites/press/list/release/release_29152.html

International Transport Forum

Blockchain and Beyond: Encoding 21st Century Transport

This report examines how advances in data science and encoding could improve transport. It investigates three linked and rapidly changing areas: First, it discusses the deployment of blockchain and other distributed ledger-based approaches, that record transactions efficiently and in a verifiable and permanent way. Secondly, the study looks at open algorithms and other alternatives to traditional data-sharing. Finally, it reviews the development of a common data syntax for encoding mobility services.

The work for this report was carried out in the context of a project initiated and funded by the International Transport Forum's Corporate Partnership Board (CPB). CPB projects are designed to enrich policy discussion with a business perspective. Led by the ITF, work is carried out in a collaborative fashion in working groups consisting of CPB member companies, external experts and ITF researchers.

International Transport Forum
2 rue André Pascal
75775 Paris Cedex 16
France
T +33 (0)1 45 24 97 10
F +33 (0)1 45 24 13 22
Email : itf.contact@oecd.org
Web: www.internationaltransportforum.org