Trends in Smart City infrastructures

István Gódor, Jan Höller
Ericsson Research
istvan.godor@ericsson.com

Keywords: smart city, big data, analysis, city administration

A smart city is envisaged to be an integrated open environment, which is able to collect all relevant information about city life, analyze this information and create a meaningful knowledge base. ICT is the enabling common infrastructure and together with a horizontal approach enables targeting data and information driven complex problems. We are looking forward to generic and open standards enabling development of open APIs and a participatory community that could drive the ecosystem into a next level. Naturally, the computation needs should build on the best performing big data analytics components; meanwhile combine them in a smart way to exploit the most of these open source solutions. Some examples are already developed and municipalities are able to use them today.

1. Introduction

The continuously changing landscape of workplaces and internal migration of inhabitants within and towards cities calls for adaptive urban planning and operations. To help authorities cope with this evolving environment and prepare the infrastructure of cities to deal with the challenges of the future, the best technologies for infrastructure for smart cities are being researched, and that require the inclusion of all stakeholders; citizens, enterprises and authorities. Having more than 7 billion mobile subscribers, digital city programs to connect meters, sensors or any devices rhyme to our vision about anything that benefits from being connected will actually be connected. In this Networked Society, people, knowledge, devices, and information are networked for the growth of society, life and business.

Such transformation of life and operations requires a digital representation of the real world updated in realtime, powered by embedded devices and connectivity and supported by the emerging 5G system and technologies to serve the needs of the Networked Society. Moreover, technology should be designed not to uncomfortably interfere with people's life. Nevertheless, to operate smart cities we need a strong cooperation between all the partners in cities ensuring the most benefit for our future.

2. What makes a city smart?

Smart devices appear more and more in our daily life, even at great celebrations like the initiation of popes. As usual in the past, both Pope Benedict and Pope Francis attracted tens of thousands people to the Saint Peter's square in the Vatican and billions of people over the television broadcasts even with the help of the Internet. *Figure 1* shows groups of people looking forward to the result of the pope elections and the first talk of the new pope in 2005 and 2013, respectively.

The impact of mobile device technology and everywhere viable connectivity is clearly evident by the massive update of smartphones and tablets in this simple illustration.

However, having such applications available for public does not make cities smarter automatically. That is, using the available technologies and having an infrastructure in operation are a must for smart cities, but not enough to be a sufficient criterion. Moreover, there are many definitions for smart cities, and applications span a very wide range, e.g. exemplified by the use cases and scenarios we were part of developing in the EU FP7 CityPulse project [1].

A smart city has to be well operated and must be a place where everything is done in cooperation with all the interested parties including the municipality, the in-

*Figure 1.
Pope Benedict the XVIth succeeded John Paul the IInd in 2005 (left) and Pope Francis succeeded Benedict the XVIth in 2013 (right)*

[source: AP, © L. Bruno, M. Sonn]

dustries, enterprises and citizens. For this to happen, a smart city must have integrated infrastructures and processes across segments, and information and communication technology (ICT) is the enabling common infrastructure that can provide this.

3. Towards billions of connected thigs

What is needed for a smart city in particular? First of all, the availability of proper and timely information is a must for operating the cities and finding the synergies between the traditionally vertically completely separated departments of the city administrations (see actual examples in Section 4). In order to have the information available all the components of the city operation should be connected in one way or another.

3.1. Transformation examples from industry segments

Nevertheless, there are some important steps ahead of us. On the one hand, we should make these components of city life to be able to cooperate with each other. On the other hand, we should understand how these components will change to support the cooperation. And these changes bring the transformation of many aspects of our life. Let us give some examples illustrated by *Figure 2*.

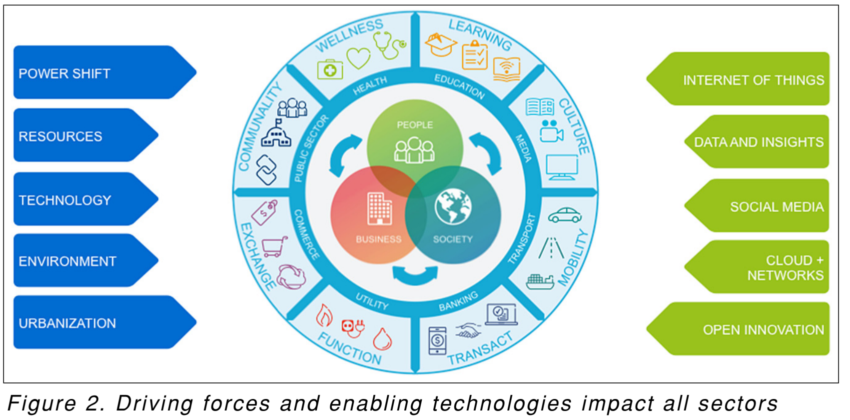

Figure 2. Driving forces and enabling technologies impact all sectors
[source: Ericsson]

The major connecting points of our life to the cities are the utility services, be it energy, transportation or other.

Energy and smart grids were one of the first targets of smart city activities. Of course the required developments in the energy industry are not limited to the technologies serving the energy distribution, but require an always-on communications, monitoring services, analytics solutions and management system beyond the capabilities of simply connecting assets of the energy infrastructure.

Regarding renewable and green energy, even in the administration area of the European Union, it is a great challenge to modernize the infrastructure that is able to cope with varying power consumption and highly distributed generation requiring adaptive feed-in management. As a consequence, individuals cannot sell their extra energy even to their direct neighbors simply because the lack of control of in- and out-flow of energy. A first step to achieve this latter is the introduction of smart meters. *"The EU aims to replace at least 80% of electricity meters with smart meters by 2020 wherever it is cost-effective to do so. This smart metering and smart grids rollout can reduce emissions in the EU by up to 9% and annual household energy consumption by similar amounts."* [2]

Another primary target example is transport and automotive, where enabling intelligent transport systems are manifold. Integration of various transport means for both people and goods are key in multi-modal transportation. We have also witnessed an increasing development of vehicle technologies in the last few years. Driver assistance systems are already available in the premium segment of cars, focusing on lane keeping/following systems, intelligent cruise control systems that are a great help on motorways, one way traffic city drive or traffic jams. Fully autonomous driving in dense areas is foreseen as highly assisted task by road side infrastructures and being heavily disputed in standardization bodies involving car manufacturers and telco companies, as well. As a future advantage for cities, the traffic will be much better controlled, and adaptive to traffic load, which can be much better redistributed in case of accidents or traffic blockage reducing the throughput of given road segments within the city. On the industry side this can mean a better tracking of goods, loading and unloading ships at docks.

Without going into detail, it is still worth to mention also emergency situation support that requires necessary tools for ad hoc adaptive measures in unpredictable events. For example orchestration and appropriate information to rescue teams, police and fire brigades to be able to more efficiently prepared to emergency situations and response to them like in case of natural disasters, accidents or panic at sport events or music concerts without bad interference with peoples' normal life.

Naturally all such examples need integrated environments and urban solutions in a sustainable way.

3.2. Be horizontal and go mainstream

What is needed to create such an integrated environment across sectors? The first step will be the digital representation of the real world for all those things, places or anything that are relevant in the physical world. Such a representation is needed to support monitoring and controlling objects, and give identities to everything.

The underlying foundation includes sensors, actuators and communications, however, the ultimate goal is to let everything be able to interact with each other, and make them be valuable part of the whole ecosystem *(see Figure 3)*.

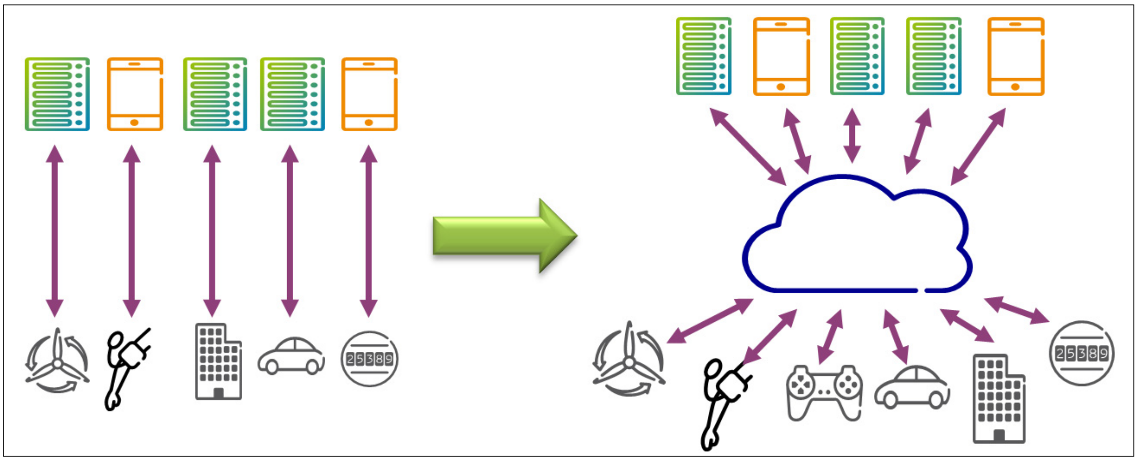

Figure 3. Illustration of today's vertical approach (left) and tomorrow's horizontal approach (right)
[source: Ericsson]

However, such ecosystems are rather limited today. Most of the solutions follow a vertical approach by focusing on a given single problem. In many cases, they are solved by specific and proprietary technologies, IT is kept in house by companies and hardly anything is shared with others. This is not the way how a smart city should be built up.

In cities we need a horizontal approach in order to be able to target data and information driven complex problems and even enable cross-domain solutions. We are looking forward to generic and open standards enabling development of open application programming interfaces (APIs) and a participatory community that could drive the ecosystem into a next level. As it was mentioned above, the strong cooperation between parties could give the most benefits to the whole city.

One of the most important things in such a horizontal system is the creation of a digital data and information marketplace which can be a part of the cloud solution given in the middle right of *Figure 3*. The broker functionality connects the users and controls how the data can be exchanged between the applications. The broker is responsible for agreeing with the participants about which sensor data or information is shared, what format is used, how often the data is reported and what incentives the participants get. In such a system the negotiation is limited to a single broker and not via a separate functionality dedicated for all applications individually. Of course, the application logic can be dedicated for the given application, but much wider selection of correlations and analytics can be used from the enriched data input pool.

In order to create such broker functionality, the interoperability of devices should be improved. It can be achieved by going mainstream.

3.3. Add intelligence to data

Once all data sources are connected together in an interoperable way, let us summarize how to get from the raw data to the intelligence. There are several steps from the atomic services from sensors and actuators, via resource abstraction, the proper semantic annotation, analytics and knowledge representation. Here we would like to summarize maybe one of the most interesting tasks that is how to close this intelligence creation loop with cloud compute and control solutions.

Why to choose cloud control? Because it has quite some advantages and it is an essential component a cooperative environments and horizontal solutions that are envisaged for smart cities. Especially, if we think on a city level, we are certainly talking about big data to be handled, which call for cloud technology.

There are three main benefits we can get from cloud control:

1. The first benefit is the possible access to a richer context that a cloud compute and control should support. This richer context might include: static information sources, periodic reports, push or poll data sources, real-time data sources, and learned historical patterns, as well.
2. The second benefit is the collective knowledge behind. Since more participants are envisaged, so more knowledge can be shared between the parties of the smart city ecosystem. We can learn patterns from all, find the correlation between them. And finally create data-driven models much better describing the world out there than anyone can calculate in a dark room of the city administration.
3. Finally, cloud simply removes the integration barrier across multiple stakeholders by relying on the same infrastructure.

In summary, the benefits are rich context, collective knowledge and easier management.

In order to enable and utilize those nice benefits given above, one needs a cloud compute platform. Naturally, such a platform builds on the best performing big data analytics components; meanwhile combining them in a smart way to exploit the most of these open source solutions. Even if it is surprising to see, the best solutions for big data analytics are coming from open source communities today.

These communities are far beyond young engineers and student coders, but supported by the world's largest companies contributing either directly to these projects or via their academic partners. The companies include IBM, SAS, SAP and many others including all major telecom vendors, as well.

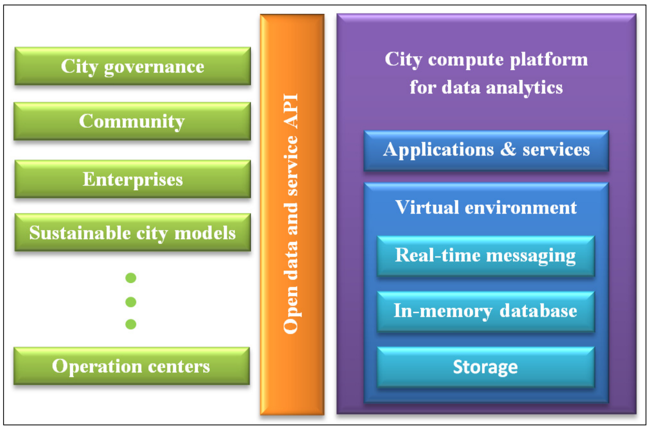

Figure 4. Example of a city compute platform

An example of a city compute platform is illustrated in *Figure 4* where all the applications and services requested by the city can be accessed via an open data and service API.

The platform itself can be realized in a cloud solution, where container technologies like Docker [3] can provide a lightweight virtual environment for the users. Such technologies group and isolate processes and resources from the host and any other containers. Technologies like Yarn [4] allow splitting up the functionalities of resource management and job scheduling/monitoring into separate daemons via a container in a container concept. By this, "application containers" can include special libraries needed by the application, and they can have different versions of Perl, Python, and even Java.

When running the actual applications and services, the industry moves towards Spark [5] providing data parallelism and fault-tolerance. Naturally, Spark has built in support for the above cited languages being able to handle both static and streaming data. Even streaming SQL can be handled by Esper in Spark. Via these languages the users can handle files, offline and in-memory databases, and streams, as well. Nevertheless, Spark is included in most Apache Hadoop distributions [6].

Beyond the software environment, the virtual environment should have efficient storage solution. For example, Hadoop Distributed File System (HDFS) [7] is the de-facto big data storage solution today. Complementary to file based data access, in-memory databases are often needed to provide a high-speed data lookup or correlation solution. Among some options, Redis [8] is often used because of its fault tolerant and scalable properties. Since there are many subscribers to the data coming into or calculated by the system, a messaging solution is needed to provide a flexible data access. There are several alternatives for messaging including ActiveMQ [9], ZeroMQ [10] and Kafka [11] all having different strengths; while no ultimate solution is applied in the industry.

In such an environment, application and service developers can work independently meanwhile having efficient and controlled way of access to all data available in the city environment.

Thereby, such a compute platform can serve the computational needs of a smart city. Nevertheless, a city architecture is not just a data analytics platform, but much more. Just consider the possible technologies providing the connectivity for the included parties (like data sensors, users, decision makers, etc.), where 5G technology will play a central role. However, going into details of 5G is beyond the scope of the paper.

4. What is a smart city? – via examples

The smart city is an integrated open environment, which provides analytics, knowledge and automation for infrastructures, things and people, i.e. all involved stakeholders in a city environment, see *Figure 5* below.

In this section, we present some examples well illustrating the data and the intelligence created from the data. The CityPulse project (an EU FP7 project) has collected hundred and one use cases relevant for city life,

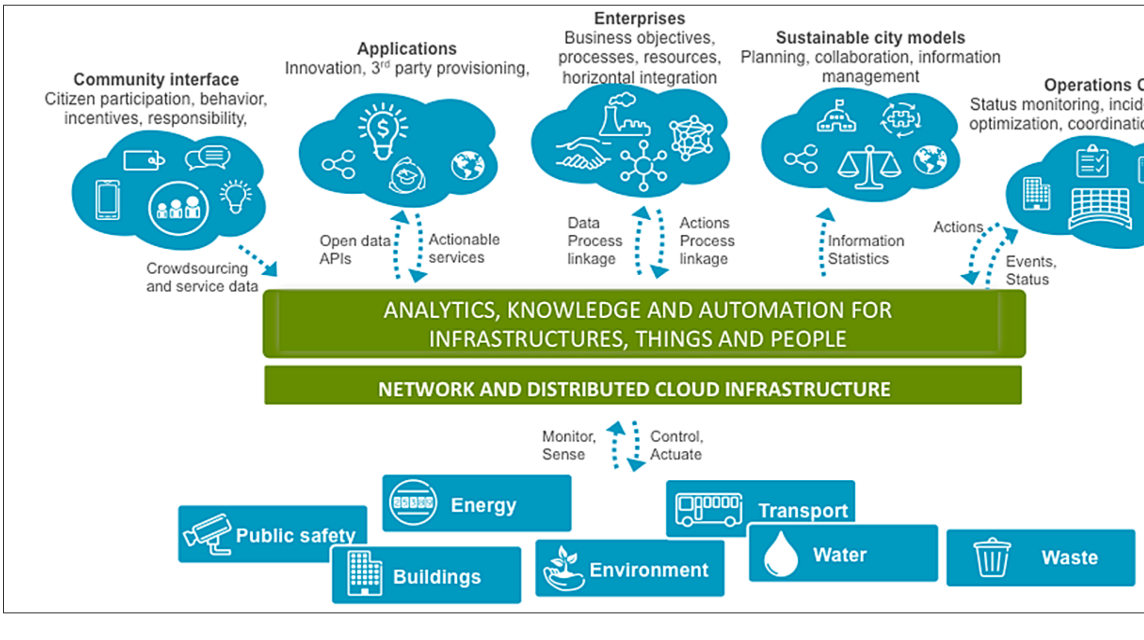

Figure 5. Stakeholders in a city environment
[source: Ericsson]

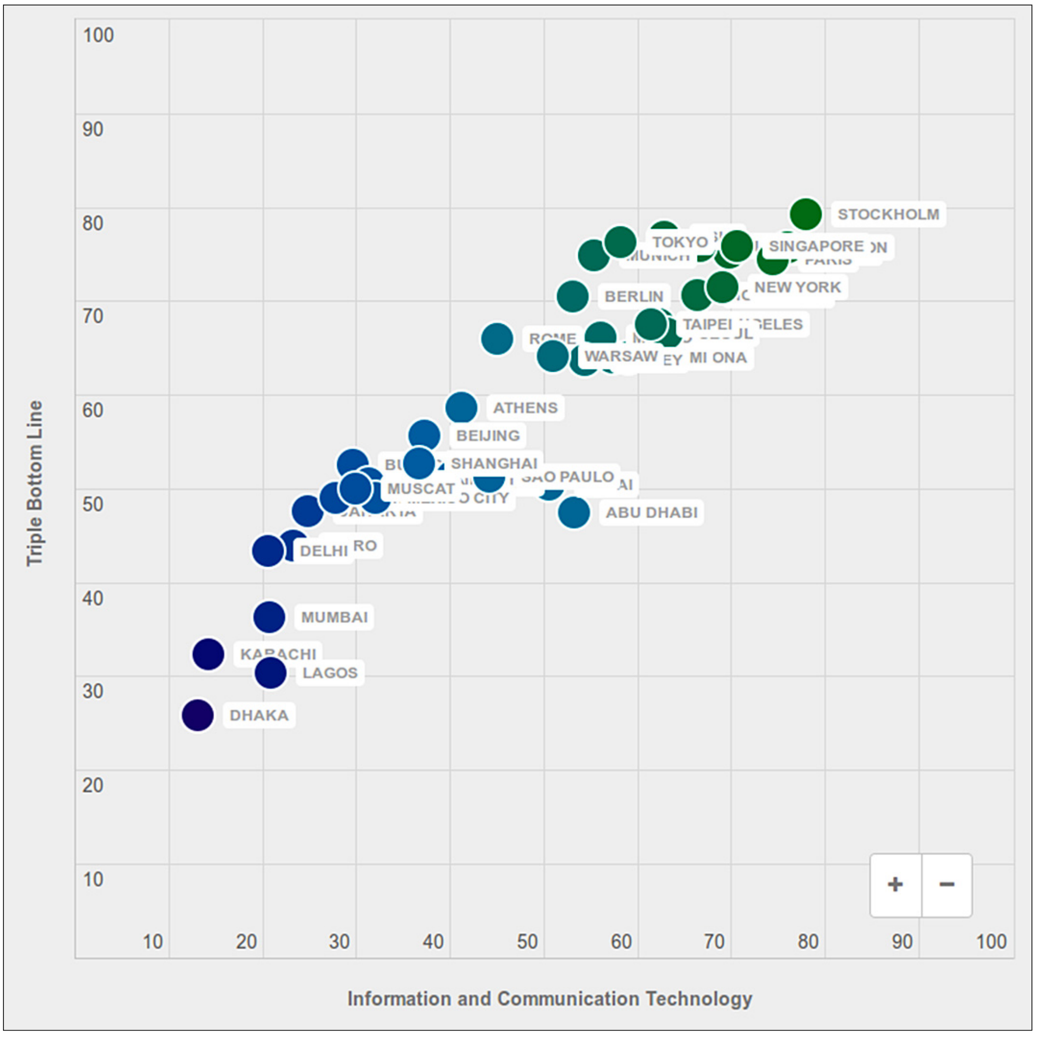

Figure 6. ICT maturity versus Triple Bottom Line development [12]

and classified them based on their relevance, input data type and decision support among others [1]. Such use cases covers multi-modal public transport, parking management, public lighting, pollution monitoring, adaptive traffic routing and many others from which we have selected five aspects and particular examples being already developed and municipalities are able to use them:

1. Be able to measure your city and compare with others.
2. Know what is happening in your city and understand its structure.
3. Make city data open and share them with all interested parties.
4. Encourage cooperation of parties.
5. Utilize the cooperation.

4.1. Understand the city

A city has to measure its status and compare it with other cities around the world. In Ericsson's Network Society City Index, we investigate how ICT maturity of cities is connected to their Triple Bottom Line development including social, environmental and financial maturity [12].

As you can see in *Figure 6,* the Networked Society City Index shows a strong correlation between ICT maturity and Triple Bottom Line development. Overall, Stockholm ranks as number one in ICT maturity dimension, while also ranking well for Triple Bottom Line. Other cities that show high ICT maturity include London, New York, Paris and Singapore.

4.2. Follow the real-time structure of the city

Once the city understands its status and position in the world, the next step is to maintain a real-time knowledge about what happening in their city.

For example, a structure of the city should be known in a more up-to-date form than a 10-yearly census poll of citizens. Our study has shown that one can understand the structure of city from the spatio-temporal variation of telecom traffic activity and thereby can have a practically real time follow-up about how the evolution of the city takes place actually [13].

Figure 7 illustrates how the structure of city is reflected in mobile traffic activity [14]. People use various applications depending on the time and their location. *"In residential areas, many people start their day by checking weather forecasts, news or social network apps, while in the evening they watch video, play games or communicate, including extensive use of SMS. Recreational areas like amusement parks, golf courses and public parks show similar usage patterns to residential areas, but with greatly increased traffic activity over the weekends. In business areas, people call and*

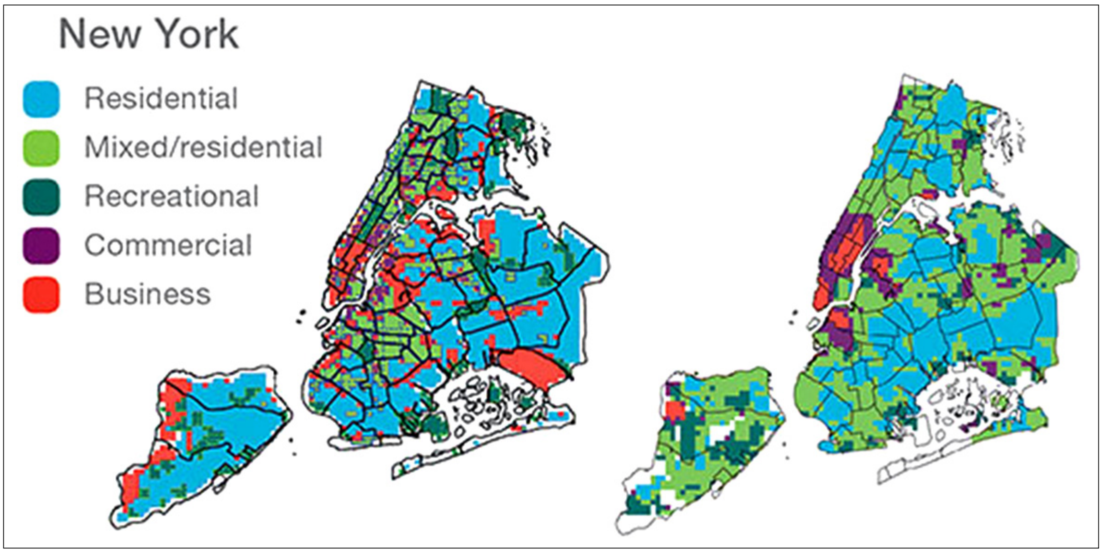

Figure 7. Comparing land use and mobile traffic patterns: official census-based land use map (left) versus traffic-based land use map (right) [14]

[source: NYC Department of City Planning]

text their clients and business associates with high intensity, mainly during work hours. Commercial areas show similarities with business areas, but with more activity during early evening." [14]

4.3. Make city data open

Similar analysis can be done for all aspects of city life including energy or water usage of different districts of the city, commutation patterns within the city, how easy to access education, healthcare or other services of the city. Naturally the administration of cities has limited resources to analyze all the possibilities and aspects of how to make their city even better based on the knowledge the city can collect about its structure and actual operation. Thereby it is essential to involve all the interested parties into this analysis via letting this information be an open knowledge.

For example, we were given an unprecedented opportunity from operators around the world to present telecom activity of cities for almost a year in an open web tool [15]. You can check Los Angeles, New York, London and Hong Kong and their various aspects illustrated in *Figure 8* like:

– Density maps as a user defined function of population and telecom activity in arbitrary combination;
– You can check the land use maps of the cities based on user activity patterns,
– And you can check the profiles of the districts of the cities.

4.4. Let us cooperate

The operation of smart cities requires strong cooperation between partners of the cities. To highlight the importance of cooperation between partners of the cities, let us present a banal example of garbage trucks.

Why do these trucks collect only the garbage and do nothing more, meanwhile they are going around the whole city day-by-day? They should check the street lights, the potholes or in general the status of the environment maybe reported via pull from the truck towards all the sensors deployed along the roads.

These connection points should be identified in more and more segments of city life, where citizens can be a great supporter of the administration. Then the administration can encourage the partners to cooperate. Many

Figure 8.
Snapshots from mobile activity around the world available at manycities.org*: density (left), land use (middle) and district profiles (right) [15]*

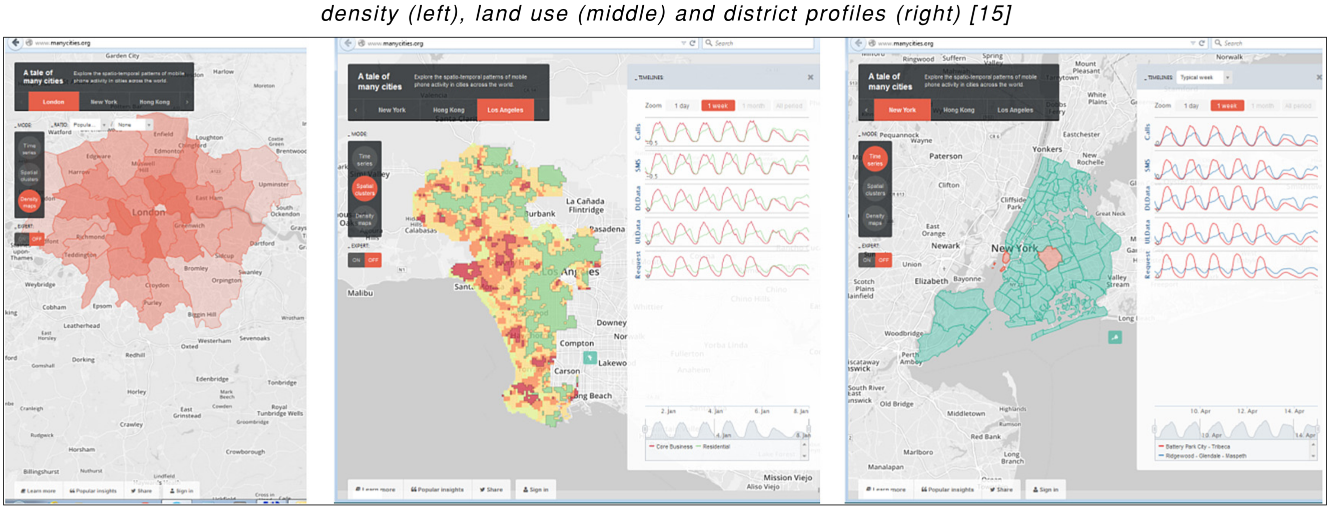

cases these partners are companies own and operated by the city, so saving money could be a great incentive to utilize these possibilities. Another aspect of cooperation is vital in public safety and emergency response tasks of the cities. For example, a dedicated coordination center could be vital in areas having many bushfires, like in Australia, where such a center and related functionalities are being developed together with Telstra (Australia's largest telecom company).

5. Summary

Smart city is envisaged to be an integrated open environment, which is able to collect all relevant information about city life, analyze this information and create a meaningful knowledge base that can be updated even in real-time where it is needed. Based on these, the city can create automation and control system for its infrastructure thereby vitalizing both things and people.

The city is the right entity to drive the transformation and utilize the advantages of a horizontal approach in the structure of city operation and city life in general. For that, cities can utilize the technologies developed in the IT world, can utilize the big data technologies including cloud compute and control solutions. In order to start up these technologies, the city has to be able to measure and compare its status even in real time, and create meaningful data driven city models.

To ensure this, the most important step is to be open for this trend and cooperate.

References

[1] EU FP7 CityPulse project
http://www.ict-citypulse.eu/page/

[2] European Commission, Smart grids and meters,
https://ec.europa.eu/energy/en/topics/markets-and-consumers/smart-grids-and-meters

[3] Docker,
https://www.docker.com/what-docker

[4] Apache Hadoop YARN,
https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html

[5] Apache Spark,
http://spark.apache.org/

[6] Apache Hadoop,
http://hadoop.apache.org/

[7] Hadoop Distributed File System (HDFS),
https://hadoop.apache.org/docs/r1.2.1/hdfs_user_guide.html#Overview

[8] Redis
http://redis.io/

[9] ActiveMQ
http://activemq.apache.org/

[10] ZeroMQ
http://zeromq.org/

[11] Apache Kafka
http://kafka.apache.org/

[12] Ericsson's Network Society City Index, 2014,
http://www.ericsson.com/thinkingahead/networked_society/city-life/city-index/graph

[13] S. Grauwin, S. Sobolevsky, S. Moritz, I. Gódor, C. Ratti: Towards a comparative science of cities: using mobile traffic records in New York, London and Hong Kong, Computational Approaches for Urban Environments, Springer, pp.363–387, 2015

[14] I. Gódor, Zs. Kallus, D. Kondor, S. Grauwin, S. Sobolevsky: The Signatures of City Life, Ericsson Mobility Report, pp.28–29, November 2014

[15] MIT SENSEable City Lab and Ericsson project:
http://senseable.mit.edu/manycities and
http://www.manycities.org

Authors

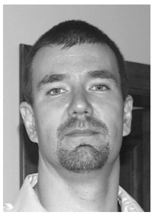

ISTVÁN GÓDOR is a senior researcher at Ericsson Research, Traffic Analysis and Network Performance Laboratory of Ericsson Hungary. He has received MSc and PhD degrees from Budapest University of Technology and Economics. He is a member of the IEEE and a member of Public Body of Hungarian Academy of Sciences. He has been awarded the 2014 IEEE Communications Society Fred W. Ellersick Prize. His research interests include network design, combinatorial optimization, cross-layer optimization, self-organizing networks, energy efficiency, traffic analysis & modeling and smart cities.

JAN HÖLLER is a Fellow at Ericsson Research with the responsibility to define and drive the Ericsson IoT research strategy and agenda. He is also contributing to the overall IoT strategy of Ericsson and in particular on IoT technology strategies including standards and industry alliances. Jan was an early promoter of Internet of Things and M2M at Ericsson by establishing the research activities in the area over a decade ago, and has over the years made significant contributions to Ericsson's early portfolio and standardization strategies on both the Group and Business Unit levels. He has been active in a number of international research programmes on the Internet of Things, and is co-author of the book "From Machine-to-Machine to the Internet of Things: Introduction to a New Age of Intelligence". He is a frequent speaker at international events ranging from academic conferences to industry shows. Jan also serves as secretary on the Board of Directors at the IP for Smart Objects Alliance, the industry's first IoT alliance founded in 2008.