19

Intelligent Infrastructures

DIMITRIS PAPANIKOLAOU[1]
MIT

IT HAS BEEN a common misbelief that intelligence and coordinated behavior is the result of a single mind's activity. In his seminal work *The Society of Mind* Marvin Minsky described a new model to explain intelligence consisting of a distributed social network of connected agents, the behavior of whom is driven by their personal goals, beliefs, and constraints as a response to external stimuli. Minsky argued that intelligence is thus not a single-mind phenomenon, but instead the product of properly wired collective behavior. Today, current research on distributed sensor networks uses market mechanisms and basic micro-economic behavior as the means to create emergent patterns of collective intelligence for managing resource allocation among the sensor nodes. How can these ideas be used to create smarter urban environments?

SMART CITIES

As William Mitchell points out, *Smart Cities* are intelligent networked urban environments that can be more responsive to the needs of their inhabitants by providing customized services on demand, making thus a more sustainable use of their resources.[2] They can do that by collecting and analyzing data from human behavior patterns through dense neural networks of sensors and microcontrollers, evaluating the data based on some higher-level goals, and providing the output back to the users to affect their behavior, closing thus a big control feedback loop. For example, smart power grids can sense consumption patterns from smart meters, correlate them with current production patterns to forecast the system tendency, and dynamically modify electricity tariffs between buyers and sellers for valley filling and peak shaving. Likewise, networks of smart domestic appliances can monitor consumption and production patterns, and determine the best time for time-elastic operation such as defrosting, washing, etc. Or, smart vehicle sharing systems can estimate fleet distribution asymmetry and incentivize users accordingly to rebalance the fleet.

Controlled behavior through feedback is not a new idea; Norbert Wiener formalized it in *Cybernetics*. For example, when you touch with your finger a hot surface your sensor cells will react, sending pain signals to your brain cells, which will evaluate their magnitude, and send back accordingly orders to your hand's muscle actuator cells to react. Nowhere in this feedback loop exists one single cell yet all cells collaborate to create what seems to be from outside a cognitive behavior. What is new though here is the idea that this cognitive behavior can be collaboratively achieved by distributed urban infrastructure systems to create responsive environments that behave as single organisms. But if this is the case, then where does coordinated behavior come from and how can it be achieved? Most of the required technology is here; it is time now to think how to put the pieces together to make them work.

Three areas of current research shape the emerging field of smart environments. First are *ad hoc* distributed networks of sensors, microcontrollers, and actuators that can sense, compute, control, and communicate; second are human–computer interaction interfaces that allow end-users to actively engage in these networks through portable or embedded smart devices, formulating thus human–machine ecosystems; third are Mechanism Design and Computational Economics for creating policies to coordinate goal-seeking collective behavior and understanding stability of these ecosystems. In this chapter we present the concept of collective intelligence in urban infrastructures and discuss some of their key theoretical principles, technologies, and questions.

URBAN INFRASTRUCTURES

Urban infrastructure networks are the nerves and veins of cities that carry information and vital resources to and from their functional units to work. As Mitchell points out, cities once consisted of mere organs, skeleton and skin in which inflow of resources and outflow of waste took place through man and animal muscle power. In the industrial era cities developed extensive mechanized supply systems of hydraulic pumps, fuel-burning motors, and power grids to provide energy and resources and remove waste, which were manually monitored and controlled. Today, in the digital age cities are developing electronic nervous systems that digitally monitor and control the organs and supply systems.[3] The next step is to

develop the collaborative patterns of intelligence that will allow them to coordinate sophisticated behavior.

Urban infrastructure networks consist of *transmission* links and *storage* nodes that reallocate scarce resources or commodities between nodes to provide service to the beneficiaries of the nodes. Transportation networks reallocate vehicles from origins to destinations to provide users with mobility; hydraulic networks pump in fresh water from ponds to households while pumping out wastewater from households to processing plants to vitalize inhabitants; electrical grids transport energy from power plants to households to provide users with power to work; data cables send information bits to and from computer registries and logic gates to allow end-users to communicate. Transmission through links is actuated by motors, pumps, or electric power sources, while storage in nodes is held by parking spaces, reservoirs, batteries, or computer memory registries.

LEVEL OF SERVICE AND CAPACITY LIMITATIONS

Citizens don't need infrastructures; instead they want the services infrastructures provide when they need them. Since demand patterns are often unpredictable, and different citizens need certain services at different times and locations, the efficiency of the infrastructure system in delivering the service relies on its ability to rapidly adapt to those spatiotemporally changing patterns. However, due to limits in the capacity of both transmission and storage technologies there will always be trade-offs as demand grows; an increase in demand for a service would lead to decreasing marginal gains from that service to people as infrastructure limitations kick in. For example, limits in transmission bandwidth bring congestion trade-offs: more cars in the streets will transport more commuters but throughput speed decreases as streets get more congested.

Similarly, limits in storage capacity bring overflow trade-offs: since land is scarce, the average parking space per user diminishes as more users buy vehicles, taking up vital public urban space. In fact, urban economists know well that there is no such thing as sufficient capacity for an infrastructure system: the larger it gets, the more its demand grows such that a new saturation level occurs, inhibiting its further growth. For example, constructing wider streets increases urban development, attracting even more vehicles until a new traffic congestion level comes that eventually restrains further development. Therefore, creating self-sufficient systems is not always about endlessly increasing their infrastructure capacity, but often about inventing new organizational policies for cleverly regulating the goals and desires of their users.

CONTROL AND POLICY

Since service must always flow from some parties to others, it turns out that performance is not only a technical issue but also a political and strategic design one. Adoption of most urban infrastructures is always a trade-off between the costs of private ownership and the inefficiencies of public services. For example, on one hand the high cost of private vehicle ownership compared with their low utilization rates, and the increasing parking requirements compared with the decreasing available urban land, make private automobiles an unsustainable solution for the future of dense urban environments. In the USA the average household has nearly two vehicles, which spend around 90 percent of their time parked while they require three to five times their footprint in urban land to be able to travel from an origin to a destination.[4] On the other hand, public service networks compromise everybody to receive the service at the same time as decided by a single authority. For example, in public transit others decide for you when, where, with whom, and how you will move; schedules are inflexible, and service coverage is often driven by political motivations rather than social needs. As a consequence, many areas become privileged while other areas remain underserved. In addition, the true social costs of public transportation, hidden in complex taxation, are often much higher than their seemingly low-priced fares.

THE PRESENT

So far, most urban infrastructural systems have been hierarchical, highly regulated, and topologically centralized: a single power plant powers tens of thousands of households; a public transport carrier controls hundreds of transit lines to mobilize a city. Control arbitration is difficult, and bottlenecks, when they occur, can have catastrophic chain reaction consequences. In August 1996 a minor failure in a single power transmission line in Oregon caused, through a domino effect, a massive failure in the entire West Coast from Canada to New Mexico, interrupting service to 7.5 million people.[5] Moreover, since energy production can not just stop when you turn off your switch, and since energy storage technology is still expensive, a central power plant must constantly overproduce electric power to avoid the economic consequences of undersupply, wasting inconceivable amounts of (often) non-renewable energy. Furthermore, users are passive recipients of service deprived from any sort of control and often unaware of their own consumption patterns. The first (and current) generation of power grids does not allow bidirectional data communication; there is no way that one edge may monitor the demand or supply of the other in real time. They are not responsive nor can they sense or compute. Like a broken telephone game, information from one edge to the other will be both outdated and inaccurate due to the system's inherent latency and lack of technological tools for advanced controlling. Moreover, arbitration complexity in centralized systems also increases disproportionally as the system grows: a single telephone service provider will rewire phone calls from each individual caller to receiver, but as the number of users increases, rewiring complexity will escalate exponentially.

THE FUTURE

However, today a new generation of intelligent urban infrastructure systems is emerging that will be able to predict bottlenecks, overflows, or shortages caused by unbalanced demand patterns using sensors, and control accordingly user behavior through sophisticated incentive policies and ubiquitous communication platforms. These sustainable infrastructures will primarily consist of distributed *ad hoc* sensor networks coupled with decentralized production, consumption, and storage abilities that will exhibit self-optimizing behavior with minimum central intervention. As an example, domestically owned renewable power generators and rechargeable high-density batteries, equipped with smart sensor meters, can turn each connected household into a potential energy producer or storage provider that can competitively sell, buy, store, or consume energy resources on demand. Thus, the end-user turns from a passive service receiver to an active stakeholder, the decisions and actions of whom may affect the performance and efficiency of the overall system's ecology. The question however is how can these individual stakeholders be informed about what is collectively optimal given the absence of a central authority? Two fundamental principles can create a sustainable, cooperative, infrastructure ecosystem: *sharing* and *incentivizing*.

Sharing

Sharing, or *fractional ownership*, is a method for sharing the cost, while increasing utilization of a large resource allocation system when the aggregate demand for resources is greater than the system's capacity. Sharing empowers the end-user to freely decide when, where, or how to reallocate, produce, or consume a resource within an infrastructure system. Typically, a sharing system involves a policy that allows fractional ownership rights over the allocated resources and a network of depositories where shareholders can deposit or withdraw these resources. Banking systems with bank accounts, freight rental service networks with their trucks, and airports with their luggage carts, are only a few of the many current examples. Recently, sharing has entered public transit systems as a complementary way to provide customized personal mobility with the form of one-way bike-sharing programs, while one-way car sharing is now starting. One-way vehicle sharing systems utilize a decentralized network of parking stations and a fleet of shared vehicles. Users can pick up a vehicle from any station and drop it off at any other station (one-way trips).

Despite their great convenience sharing systems have drawbacks too. Lack of cooperation and individual selfish behavior are not able to sustain welfare in sharing systems. For example, in vehicle sharing since departures and arrivals vary randomly in stations, eventually vehicles are all ending at the stations with no demand.[6] This inventory imbalance not only decreases throughput, but it also increases trip time as drivers search for parking spaces. To maintain the level of service a vehicle-sharing system needs to constantly feed origin points with vehicles while draining destination points from occupied parking spaces. While it is possible to centrally monitor bikes and periodically redistribute them with trucks, this is clearly not a viable solution for larger vehicles such as cars. Not only is it operationally complex, but also it is expensive: either the fleet needs to be too large or the redistributions need to be too frequent. In addition, continuous redistributions keep vehicles away from the system, reducing further service capacity. As a consequence, many vehicle-sharing systems end up wasting more resources sustaining their performance than the value of the service they provide.

Incentivizing

If users in a sharing system had a common way to evaluate the impact of their actions in the performance of the system, then decision making would be straightforward, individual actions would not worsen collective output, and cooperation would emerge. Incentive-based strategies such as dynamic pricing have been successfully employed in decentralized resource allocation networks with limited capacity as a means to create feedback mechanisms to regulate demand patterns: congestion pricing zones, smart grids of renewable energy resources, eBay-style online auctions, and carbon trading programs, are just a few of the many successful examples.

To calculate payoffs, information from both the users' actions and the nodes' condition must be known. This can be done at the individual level if nodes in the network can talk to each other. Distributed networks of smart domestic appliances equipped with sensors and cheap microcontrollers can easily perform basic computation and propagate information through gossip algorithms,[7] creating ubiquitous inhabitable environments of distributed computation, a concept otherwise known as the *internet of things*.[8] Duncan Watts showed that designing such scalable *small-world* networks where any pair of nodes can be linked through a path of maximum six steps is easy.[9] For example, networks of parking stations may calculate parking prices by observing inflow and outflow of vehicles, average it with their neighbors, and communicate this information back to the users through handheld devices.[10] Information thus brings a powerful tool for decentralized control.

Calculating payoffs using sophisticated smart infrastructure would have trivial value if this information could not be perceived effectively by the end-users to affect their decisions. Creating ubiquitous communication platforms of portable mobile devices using intuitive user interfaces is the key component for closing the feedback loop in intelligent distributed infrastructure systems.

In what follows we will see how these two principles can be applied by discussing *Mobility on Demand, (MoD)*. This is a research direction for intelligent shared urban transportation systems that the Smart Cities group of the MIT Media Laboratory has been developing since 2003.

MOBILITY ON DEMAND: INTEGRATING URBAN INFRASTRUCTURE NETWORKS

MoD is an integrated proposal for the future transportation scheme of dense urban environments that consists of a decentralized network of rapid charging stations, a shared fleet of lightweight rechargeable electric vehicles, and an intelligent fleet management system.[11] MoD allows users to conveniently pick up a vehicle from any origin station and drop it off at any other destination station. Each station in MoD is equipped with an Uninterruptible Power Supply (UPS), a high-capacity battery that can slowly charge with energy from the urban electrical grid during off-peak hours and rapidly charge vehicles during daytime parking. Vehicles thus can rapidly pull energy from a charging station, transport it in their batteries, and push it back to another charging station at speeds of up to 15 minutes. Each vehicle's battery can store up to 10 KWh, enough energy to power a household for one entire day, offering thus not only a tremendous energy storage buffer for smart power grids of renewable energy resources but also an emergency energy reallocation system (Vehicle to Grid technology). Therefore, MoD consists of three synergetic urban infrastructure networks: a one-way vehicle-sharing system, the smart power grid of the charging stations, and the energy transportation network of the vehicles batteries (Figure 19.1).

To address fleet and energy distribution asymmetries MoD uses price incentives to motivate users to drive vehicles to the stations that most need them, while discouraging them to drive vehicles to stations that do not need them, and an intuitive graphical user interface to effectively communicate location-based price information. Similarly to a market economy, prices adjust to parking needs. Each station locally computes a pick-up and a drop-off price based on its inventory change rate, its available energy resources, and the price competition with its neighbor stations; these two components are then added to the standard trip fare as a negative or positive percentile discount. Therefore, some trips can be more expensive while other trips may even pay back money to the users (Figures 19.2 and 19.3). There are no trucks, nor employees involved in the fleet redistribution. MoD is essentially a self-organizing system operated by its users for its users.

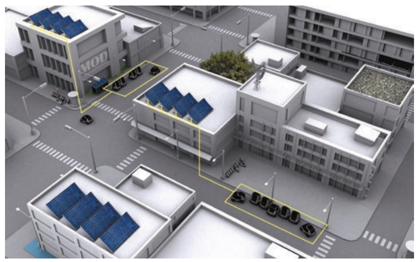

19.1 Mobility on Demand

MoD's prices are discounts off the trip fare, not fixed prices; thus, they have the same behavioral impact on trip decision making independently from the trip's length. Users can access online price information at the stations, on their handheld mobile devices, or at desktop computers with internet access. Users can either prepay during pick-up (thus locking the price) if they know their destination station in advance, or they can pay during drop-off, allowing the price landscape to change during trip time. MoD thus does not force users to use it; it does not aim to replace existing transportation options. Instead it aims to offer more options to users, creating a market competition environment among transportation options. Users that consider a trip to be expensive may simply opt out, choosing to perform the trip using another option (e.g., walking, taxi, public transit, or private automobile).

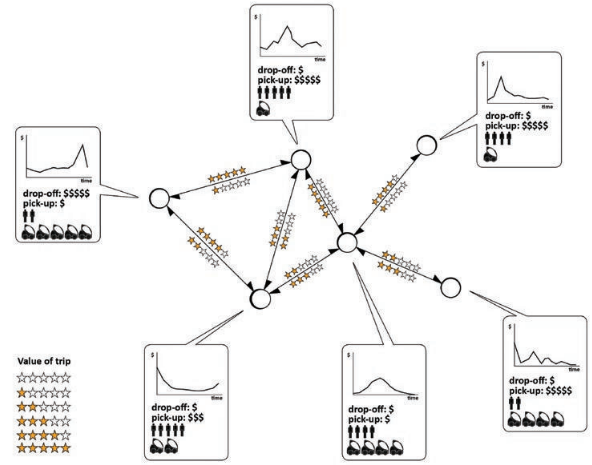

19.2 The market economy of trips
Source: William Lark, Jr. and Michael Lin, MIT Media Lab, Smart Cities

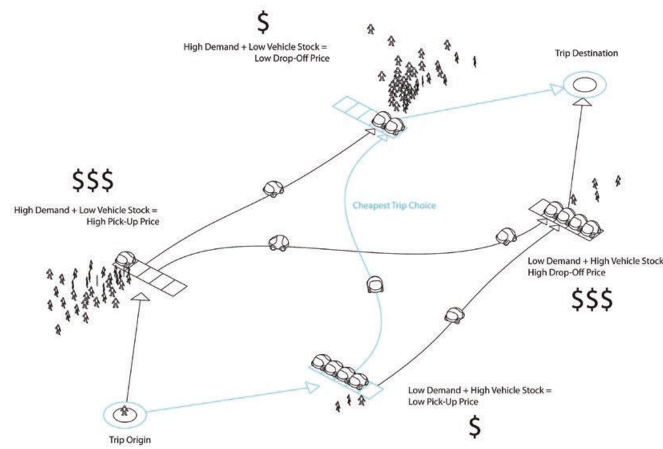

19.3 User case

LANDSCAPES OF PAY-OFFS: VISUALIZING INFORMATION

Understanding the pay-off landscape is necessary for users of MoD to bring it into a sustainable equilibrium. Smart Cities has been developing PriceScape, an intuitive web-based Graphical User Interface that uses dynamic heat-map and contour display for communicating location-based price information to users. Isometric price curves or color zones describe areas with the same parking discount rates. Like the analogy of navigating through a price landscape, climbing from valleys to hills is expensive, while descending from hills to valleys is rewarding. Traveling between locations of the same level is neutral (Figure 19.4). Users can scroll in time to see the dynamic trend of it. PriceScape is not a recommender, nor an expert system; it does not suggest to users what to do; however, it helps them to perceive their pay-offs to assist their decision-making process.

MOD ECONOMICS

How would users make decisions given this pricing context? Each urban trip is inevitably a combination of at least two options (one of which is always walking) and users will select that combination which minimizes total trip costs. As an example, consider a user traveling with a MoD system from an origin O to a final destination D (Figures 19.5 and 19.6). The user will select an in-between drop-off station Q if and only if the total time-adjusted trip cost from the origin to the drop-off station with MoD (including the parking discount) plus the time-adjusted trip cost from the drop-off station to the final destination with the public transit (e.g., bus, taxi, walking, etc.), are in sum less or equal than the original time-adjusted cost that the user would pay to travel from the origin to the final destination with MoD (Figure 19.5). Since users evaluate time differently based on their level of income, each user would select a unique drop-off station for a given price condition.

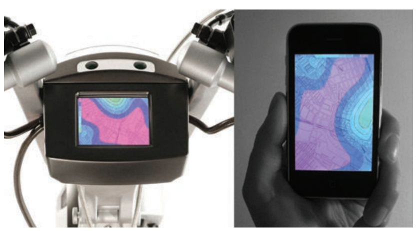

19.4 User interface in vehicle and portable devices

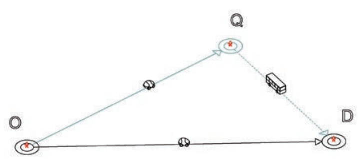

19.5 The decision-making process

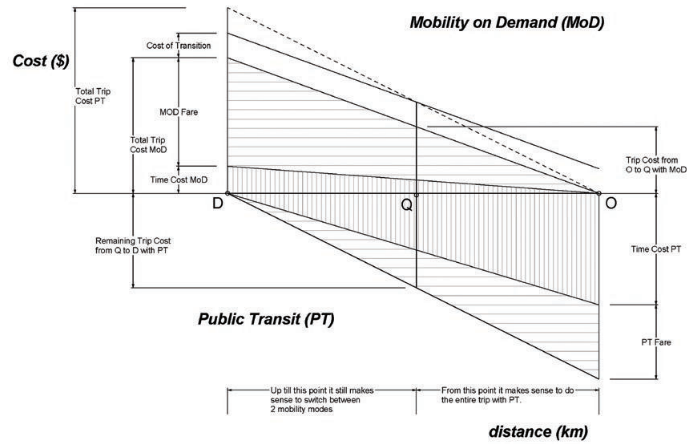

19.6 Market equilibrium between two transportation models

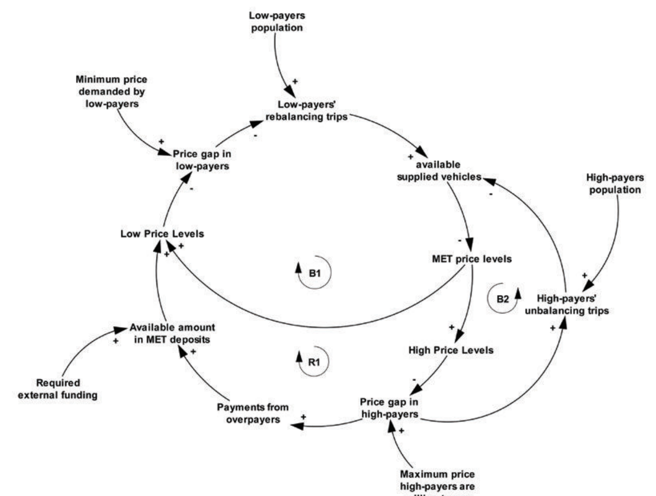

19.7 Causal-loop diagram explaining behavior of a vehicle-sharing system

MoD, as other intelligent infrastructure systems, is in fact a form of a *strategy game*. Decisions of individual price-taking users change the price landscape affecting decision making of other users. Contemporary urban economic theory[12] suggests that users with sufficient information would make decisions that minimize their time-adjusted costs, eventually bringing the system into an *equilibrium* state where no further action can be taken to increase their pay-offs. In that ideal state, the excess of money from the overpaying users plus any additional external funding should match exactly the reward demanded from the underpaying users. Moreover, the overall throughput performance of the system will depend on their price sensitivity; the higher the demand elasticity, the better the performance will be.

Income distribution greatly affects equilibrium in MoD. While the high-level prices for the high-payers are determined directly by the stations, the low-level prices that are offered to the low-payers are determined by the stations and the available pool of funds in the MoD deposits. These funds increase by the high-payers and decrease by the low-payers. This simply means that MoD cannot pay back endlessly to low-payers; it can only pay back money to the extent that this money exists in the system deposits, financed by the high-payers. To solve this issue either the system might need to be larger, or an additional external source of funding may be required. This can be provided either by increasing the standard MoD fare, or by utilizing external funding sources such as advertising etc. Figure 19.7 shows a causal-loop diagram in *System Dynamics* that graphically explains this important concept. Polarity of arrows indicates how the effect is related to the cause. Loops can be either self-reinforcing (R) or self-balancing (B). B1 and B2 balancing loops determine the equilibrium between low-payers and high-payers. However, R1 reinforcing loop may drain available funds, reducing rewards offered to low-payers, which would reduce their willingness to rebalance vehicles dragging down the system into a lower-performance equilibrium state.

DISCUSSION: WHERE DOES INTELLIGENCE COME FROM?

Intelligence is typically associated with the ability of a system to adapt to the changing conditions of its environment. As our societies become more complex, it is becoming increasingly evident that urban infrastructures capable of foreseeing their goals, understanding their needs, and reflecting back to their users with incentives for motivating their actions will be essential in creating sustainable and responsive environments to human needs. Intelligent infrastructures are in essence networked, distributed resource allocation markets consisting of two parties: those who control the stocks at the nodes, and those who control the flows at the links. It is the mutual interaction of those two parties guided by their personal goals, beliefs, and constraints that eventually determines the system's intelligence and the level of its sustainability.

But what is sustainability? For sure, it is not just about being green, or simply reducing CO_2 emissions. Sustainable infrastructures are infrastructures able to sustain themselves; those whose aggregate generated value can outweigh their aggregate operational, societal, and environmental costs. If this is the case, then the issues that we should start reasoning about in the future should focus on understanding the circumstances under which such systems can indeed become sustainable, as well as understanding what types of social equity they may bring.

NOTES

1 The author wishes to acknowledge Professor William J. Mitchell, Kent Larson, and the Smart Cities group for all the help and guidance on this work. Figure 19.2 credits: William Lark, Jr. and Michael Lin, MIT Media Lab, Smart Cities. All other figures are courtesy of the author.

2 W.J. Mitchell, *City of Bits: Space, Place, and the Infobahn* (Cambridge: MIT Press, 1996).
3 W.J. Mitchell, Casalegno F., *Connected Sustainable Cities* (Cambridge: MIT Mobile Experience Lab Publishing, 2008).
4 US Department of Transportation. *National Household Travel Survey: Summary of Travel Trends* (2001), pp. 11–15.
5 D.J. Watts, *Six Degrees: The Science of Connected Age* (New York: W.W. Norton & Company, 2004).
6 A. Kaltenbrunner, Meza R., Grivolla J., *et al.*, "Bicycle cycles and mobility patterns—Exploring and characterizing data from a community bicycle program," 2008, http://arxiv.org/abs/0810.4187.
7 P. Levis, Culler D., Patel N., *et al.*, "Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks." In: *Proceedings of the 1st USENIX/ACM Symposium on Networked Systems Design and Implementation* (NSDI, 2004).
8 N. Gershenfeld, *When Things Start to Think* (New York: Henry Holt & Company, 1999).
9 Watts, *Six Degrees.*
10 D. Papanikolaou, *Mobility on Demand: Sustainable, Self-Optimizing, Vehicle Sharing Systems.* Poster presented at 2nd International Sustainability Conference (MIT, 2010).
11 W.J. Mitchell, *Mobility on Demand: Future of Transportation in Cities* (Cambridge: MIT Media Laboratory, 2008).
12 D. DiPasquale, Wheaton W.C., *Urban Economics and Real Estate Markets* (Upper Saddle River, NJ: Prentice Hall, 1995).