mobility on demand

Future of transportation in cities

mobility on demand

Future of transportation in cities

Smart Cities
MIT Media Laboratory
June 2008

acknowledgments

William J. MItchell

ISBN xxx xxxx xxxxx

Special Thanks to

Ralph Glakenheimer

Chris Zegras

MIT Media Laboratory

William J. MItchell, Principal Investigator

Ryan Chin, ...

Mobility on Demand Text

William J. MItchell

Book design and layout

Christine Outram

Urban Design Team

Chi-Chao Chuang
Christine Outram
Claire Abrahamse
Quilian Riano
Sandra Frem

Business Model Team

Yaniv Fain
Marcus Parton
Andres Sevstuk

Vehicle Designs

RoboScooter: Michael Lin

CityCar: Franco Vairani

Vehicle Design Team

Retro Poblano
Wayne Higgins
Will Larke
Ethan Lacy
Bo Stjerne
Neri Oxman
Jamie Hu
Somnath Ray
Tom Brown
Brian Morris
Andrew Leone
Charles Guan
Christine Lin
Sung Hyuck Lee
Raphael Moyer
Arthur Petron

Additional Collaborators

Joshua Fiala
Carl Yu

mobility on demand

Mobility-on-demand systems provide stacks and racks of light electric vehicles or bicycles at closely spaced intervals throughout a city. When you want to go somewhere, you simply walk to the nearest rack, swipe a card to pick up a vehicle, drive it to the rack nearest to your destination, and drop it off.

Users of mobility-on-demand systems have the convenience and comfort of private automobiles without the associated high cost, insurance requirements, need to refuel, service and repair demands, or parking problems.

Key factors in the success of mobility-on-demand systems are the costs to users and the system latencies – that is the times needed to walk from a trip origin to a nearby stack and pick up a vehicle, to travel to a stack near the desired destination, and to drop off a vehicle and walk to the actual destination. Well-designed and well-managed mobility-on-demand systems should be able to provide more attractive combinations of costs and latencies than alternative systems such as private automobiles, taxis, and transit systems.

Management is accomplished through an innovative combination of:

1. Realtime, fine-grained mobility demand sensing;
2. active realtime management to balance vehicle (and parking space) supply and demand and meet latency targets at sustainable cost; and
3. sophisticated use of dynamic pricing for demand management. The mathematical model used for management represents the system as a network of stacks and links, with queues (maybe zero-length) of users waiting to access vehicles and of vehicles waiting to access parking spaces at stacks, and dynamically varying latencies and prices on stacks and links.

Since mobility-on-demand systems employ lightweight electric or human-powered vehicles, they are energy-efficient, carbon-minimal, and silent. They are compact, and they have very high utilization rates, so they minimize urban traffic congestion and parking space requirements.

Thus the essential parts of mobility-on-demand systems, as described in more detail below, are:

1. Specially designed vehicles;
2. Vehicle stacks and racks distributed throughout the service area;
3. ICT infrastructure for sensing and control;
4. Demand sensing and network management software;
5. Innovative electrical supply systems that utilize clean, renewable power sources and minimize transmission losses.

These elements work in combination to provide the benefits.

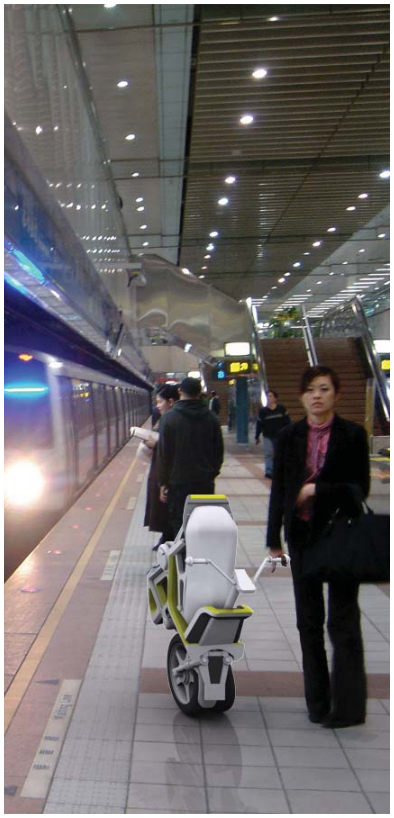

Clean, compact, energy efficient vehicles

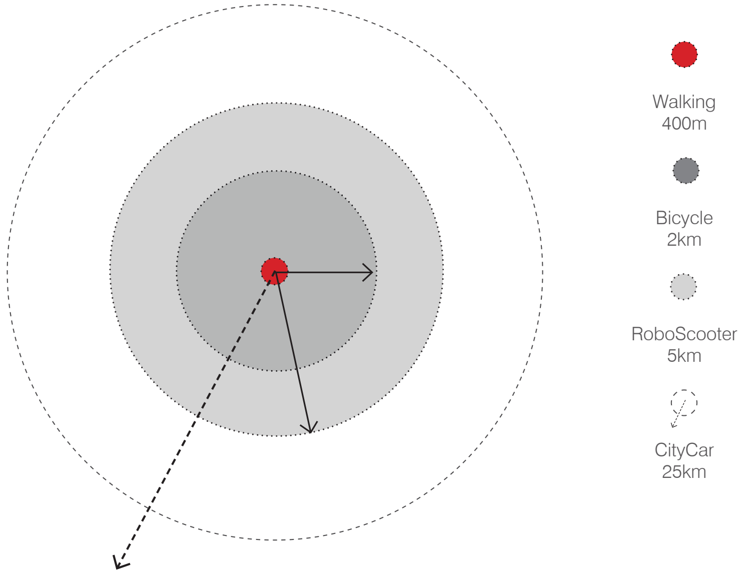

Range per mobility mode

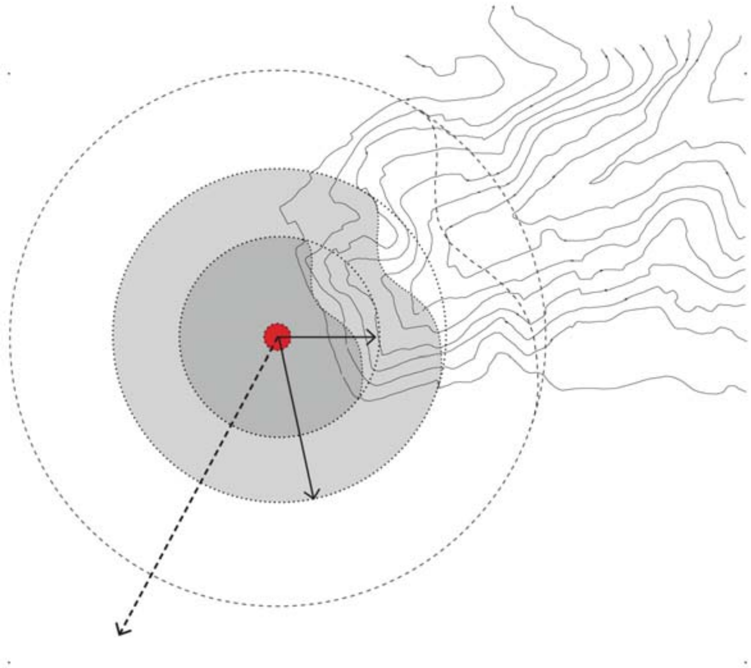

Effect of topography on mobility range

Mobility-on-demand systems may use a single vehicle type. However, a more attractive option in larger and more sophisticated systems is to employ multiple vehicle types – providing users with choices among combinations of cost, comfort, and functionality. For example, a user might choose to ride a bicycle to the supermarket, leave it there, and bring back a car to carry the bags of groceries. (Many inefficiencies in traditional urban mobility systems – for example, driving an empty SUV to the supermarket – result from the fact that vehicle types cannot be matched to trip purposes. One size must fit all.)

The CityCar, developed by the Smart Cities group at the MIT Media Laboratory, is specifically designed to meet the needs of mobility-on-demand systems .

CityCars are lightweight electric cars with in-wheel motors. They fold and stack like shopping carts at the supermarket or luggage carts at the airport, making them extremely compact and efficient in the use of urban space. They are simple and modular in their design (yet highly functional), robust, inexpensive, and easy to maintain. They recharge automatically in their parking spaces – much as electric toothbrushes recharge in their holders – so they do not need very long ranges or to carry around large numbers of batteries.

RoboScooters, developed by Smart Cities in collaboration with ITRI and SYM, are also lightweight, folding, in-wheel-motor electric vehicles. These two-wheelers are smaller, lighter, less expensive, and consume less energy than their counterpart enclosed, four-wheel cars. They also have shorter range. They are particularly suitable for use where weather conditions are good, individual transportation is the priority, and urban and economic conditions are less favorable to automobiles.

Bicycles can also be used in mobility-on-demand systems, as in the Vélo system in Paris. They may be traditional bicycles, or "smart" electrically assisted versions. These provide the lightest, cleanest vehicle options, but their use can obviously limited by terrain, weather, and range and carrying capacity demands.

Segway personal transporters have sometimes been proposed for use in mobility-on-demand systems. These may be suitable in shorter-range, lower-speed situations, and for indoor-outdoor use. (The Dutch Railways are currently exploring the possibility of Segway mobility-on-demand at railway stations.)

<< Range per mobility mode

< Effect of topography on mobility range

> Suite of vehicles used in the mobility on demand system: CityCar, RoboScooter, Bicycle, Segway

> The folding mechanism of the RoboScooter allows it to fit unobtrusively into the street

Clean, compact, energy efficient vehicles

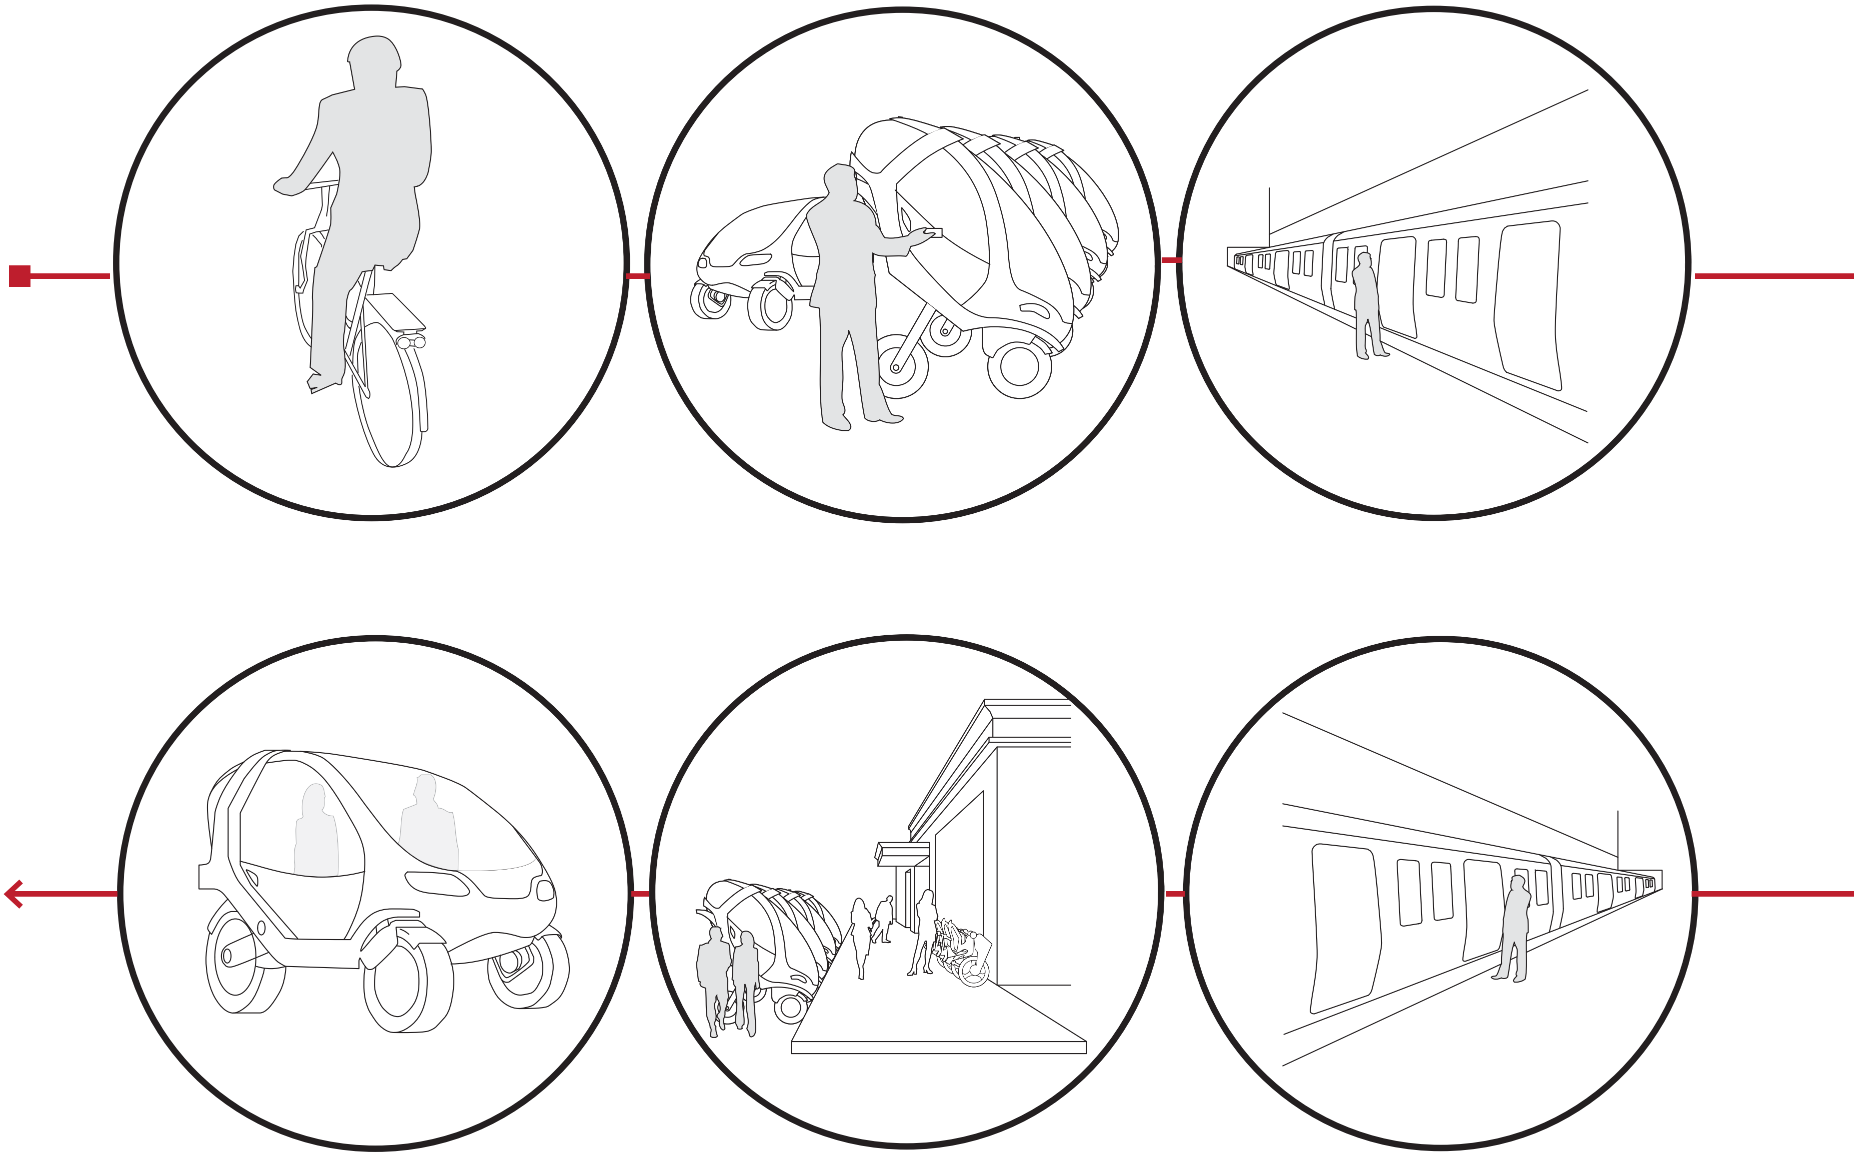

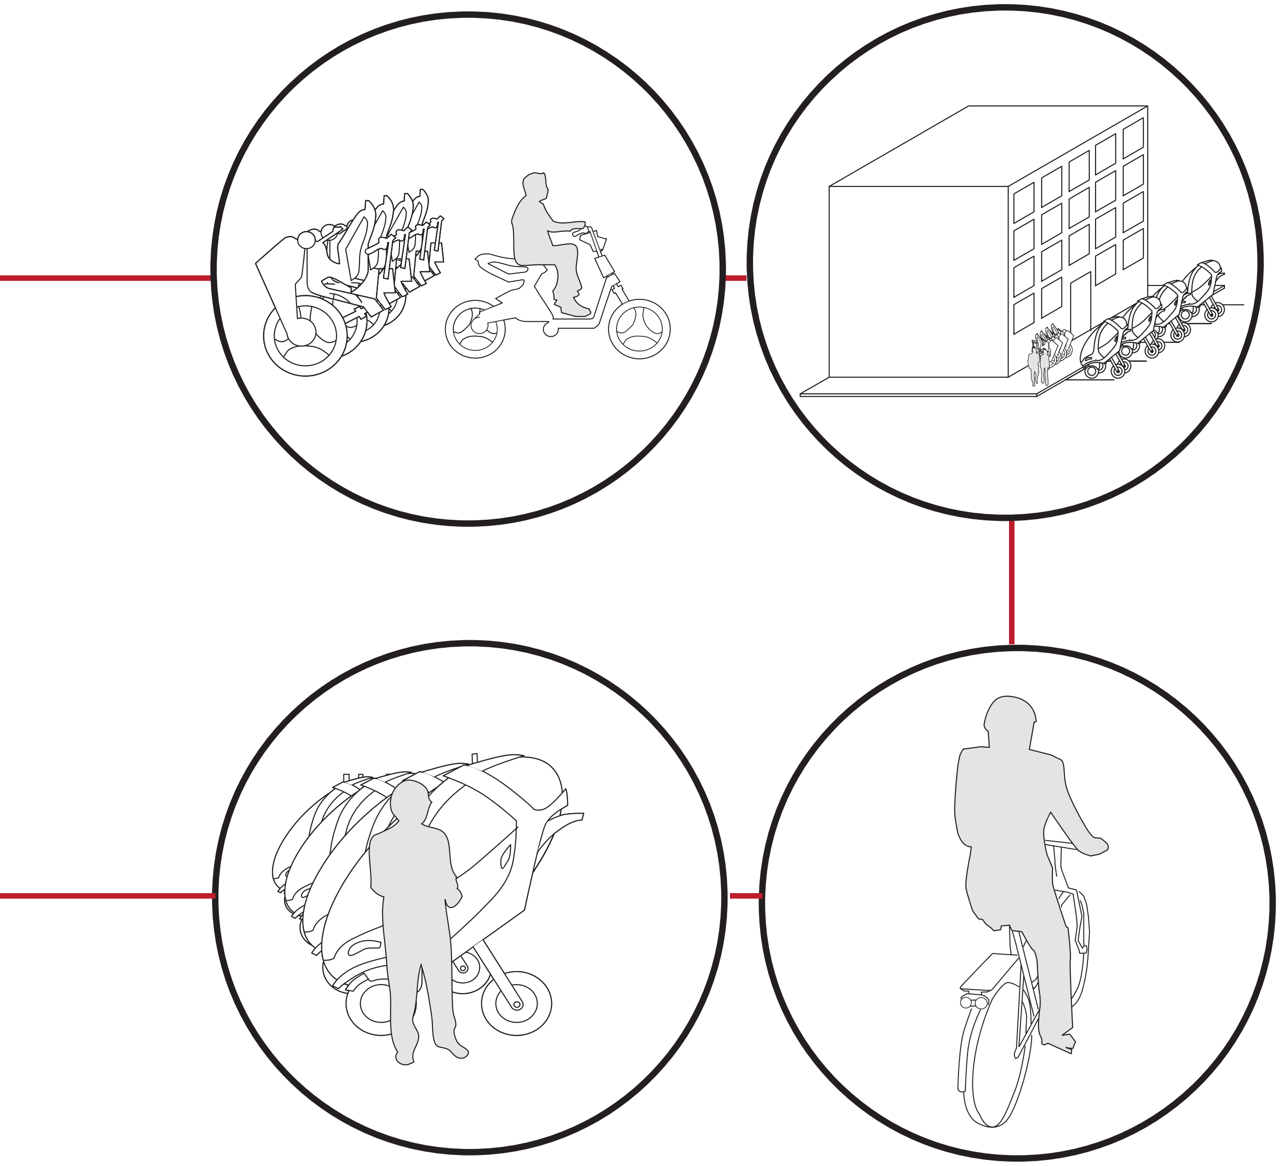

> Mobility on demand users can employ multiple vehicle types that are specifically suited to their needs and a complement to public transport.

Stacks and racks are the vehicle pickup and dropoff points in mobility-on-demand systems. (Of course, vehicles may also be parked temporarily at other locations.)

These points need to be distributed sufficiently densely around the service area to be always in close proximity to trip origins and destinations. They not only provide access to vehicles, but also enable vehicle recharging, provide vehicle security, and handle the vehicle pickup and dropoff transactions – which must be electronic, quick, and seamless. They need suitable space and street access, electric power supply, and network connectivity .

A crucial technical issue, in design of stacks and racks, is the provision of efficient, safe, convenient, and weatherproof connection between power supply and parked vehicles. The connection might either be through contact or through induction. In any case it should be automatic whenever the vehicle parks, so that there is no need for the user to "plug in" or perform any other explicit action. The idea is that users never have to think about refueling or recharging; the system simply provides charged vehicles .

These ubiquitous access points allow the system to operate in one-way rental mode, rather than two-way rental as with traditional car rental systems. Instead of relying upon users to bring vehicles back to pickup points, which simplifies management but greatly reduces the flexibility and responsiveness of the system, the operator accepts the responsibility (and reaps the rewards) of managing the distribution of vehicles in the system so that they are always available to meet demand. (Two-way rental can be regarded as a restricted special case of one-way rental, implemented by means of price incentives to return vehicles to pickup points, and managed using the same technology.)

Locations of stacks and racks will be determined by some combination of urban design considerations, availability of suitable sites, and long-term patterns of demand. Often it makes sense to combine them with existing service points, such as convenience stores, coffee shops, hotels, or bank ATMs – to the commercial benefit of both. Obviously, as well, they can usefully be placed at major origin and destination points, such as railway stations, office towers and parks, and sports and entertainment facilities .

Some stacks and racks may be small, informal, and temporary. Others may be large, permanent mobility interchange points incorporating retail and service facilities that take advantage of the traffic passing through. Larger nodes in a system may serve as vehicle cleaning and maintenance points .

Stacks and racks may be deployed incrementally as a mobility-on-demand system grows, increasing both area and density of coverage. And locations may be adjusted, over time, in response to experience of operating the system. Stacks and racks are modular, and (unlike subway stops, for example) are not necessarily locationally tied to fixed infrastructures .

< In Taipei City, the ubiquitous 7-11 network at 5 minute intervals provides an excellent opportunity for the insertion of scooter stacks

> In Florence, the scale of stacks and racks could ranges from large storage areas outside of the historic center, key mobility nodes at the traditional city gates, semi-permanent nodes that relate to the existing piazzas and portable snap on street elements that could be adjusted once the system is developed.

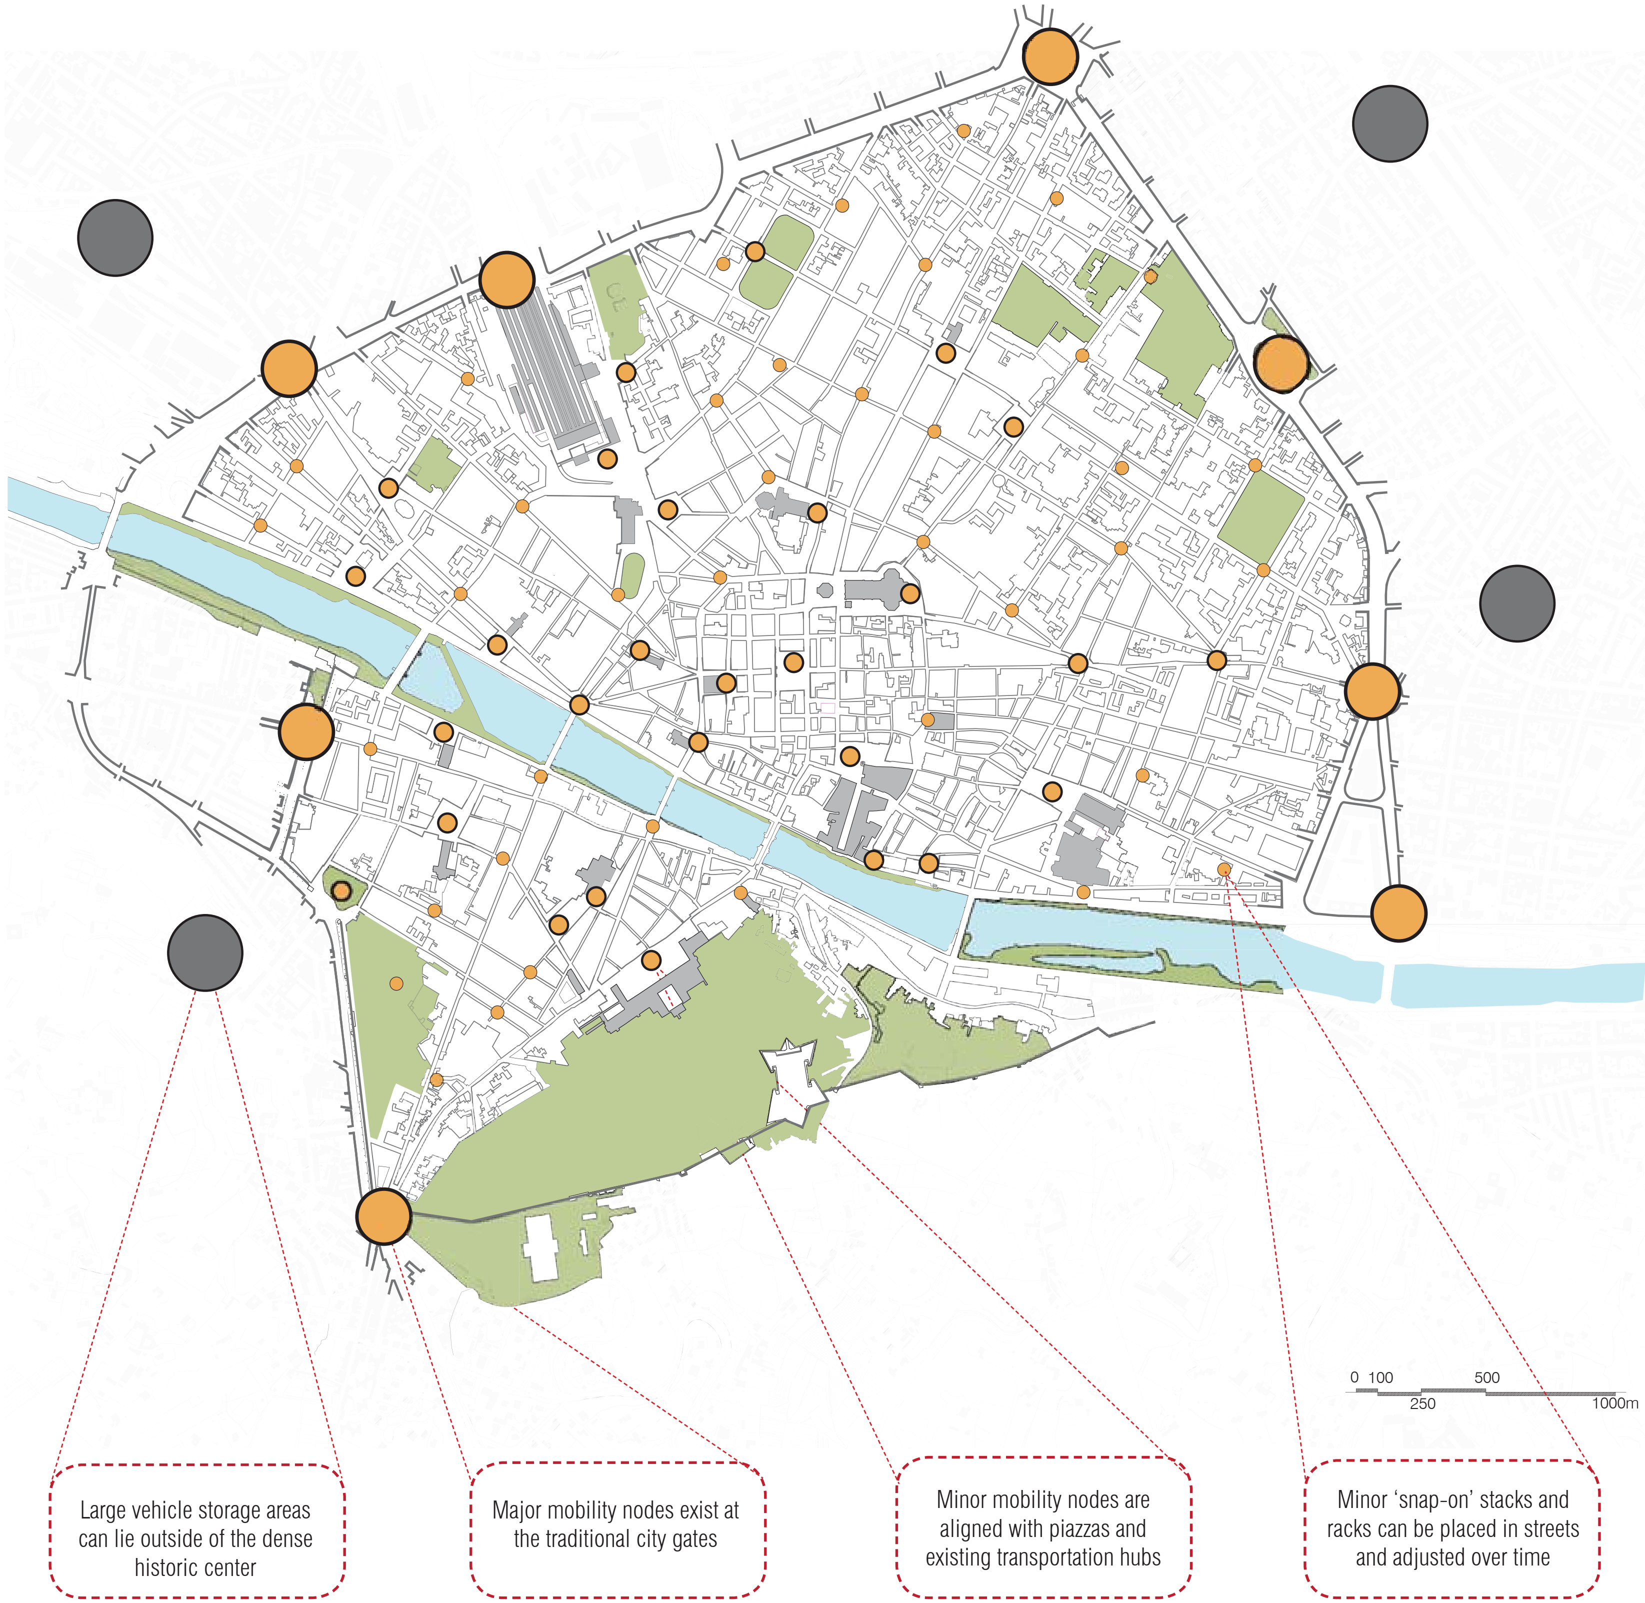
0 100 500
250 1000m
Large vehicle storage areas can lie outside of the dense historic center
Major mobility nodes exist at the traditional city gates
Minor mobility nodes are aligned with piazzas and existing transportation hubs
Minor 'snap-on' stacks and racks can be placed in streets and adjusted over time

Stacks and racks throughout the city

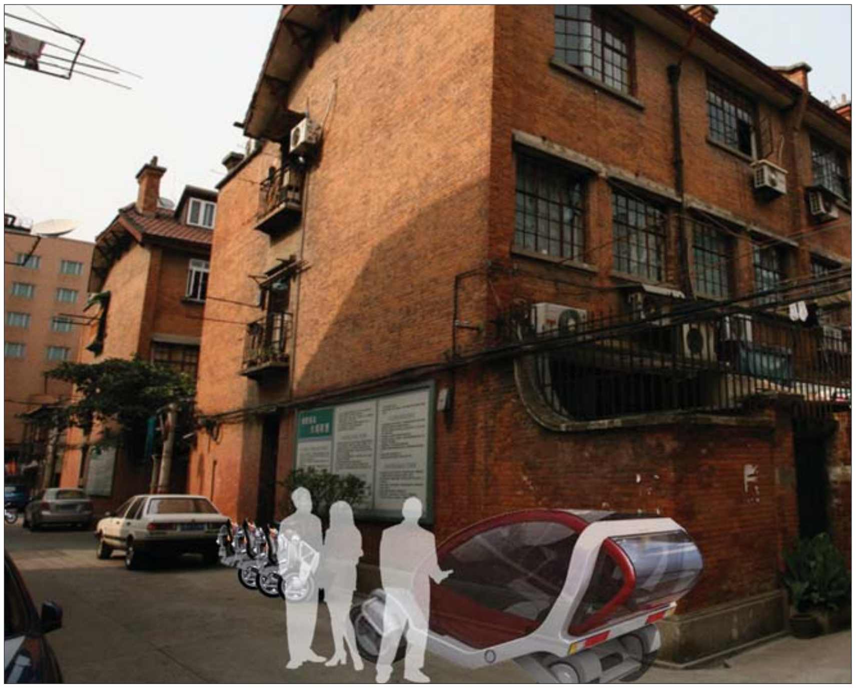

Placing mobility on demand points in the narrow streets between traditional Lilong housing in Shanghai would free up valuable public real estate that is usually consumed by parked cars

Mobility on demand outside the MIT Media Laboratory

Mobility nodes can be combined with existing traditional city gates in Florence to form new transport hubs

In San Francisco, major mobility-on-demand nodes can be combined with important civic buildings

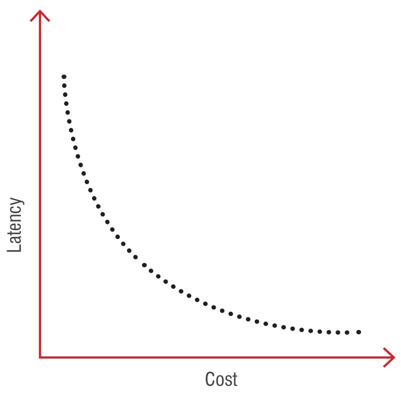

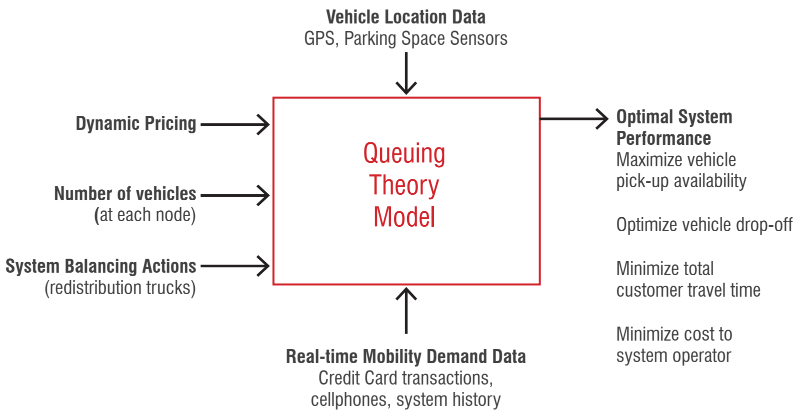

<< Relationship between trip cost and latency

< Elements in the queueing theory model

> Pickup and dropoff nodes are located throughout a cities street system creating a network of mobility-on-demand points. Total trip time can be calculated by adding latencies in the system

When mobility-on-demand pickup and dropoff points are located within a city street system they form a mobility network. The pickup/dropoff points are nodes, and the streets provide the links among them.

Each node can park some finite number of vehicles. At any moment, a node may or may not have vehicles available, and it may or may not have empty parking spaces available. Ideally, whenever a pedestrian walks to a node there is a vehicle available for pickup, and whenever a driver approaches a node there will be a vacant space to drop off the vehicle. In practice (particularly when the system is heavily loaded), pedestrians and drivers will sometimes have to queue to get access.

It is possible to gather information on the lengths of pedestrian queues at pickup points. There might be some sort of sensor system. Waiting users might "punch in" to signal that they are waiting for a vehicle at that location. Or users might employ their cellphones, as in calling for a taxi, to inform the system that they will want a vehicle at a particular location and time. (This is functionally equivalent to making a reservation.) In any case, the system operator will want to manage the system in such a fashion that it directs vehicles to pickup points where there are queues of waiting customers – just as a taxi dispatcher might.

Similarly, it is possible to gather information on the lengths of car queues waiting to park at dropoff points. (Vehicles driving along the road towards a dropoff point are implicitly queued – they don't have to be lined up at the dropoff point, waiting to get in.) Most obviously, this can be harvested from destination information that users punch into GPS navigation systems. Queue length for a dropoff point can also be inferred from the numbers of vehicles originating at other pickup points that are now in the vicinity of the dropoff point.

The network therefore forms a queuing system, somewhat analogous to a packet-switching network such as the Internet. Vehicles travel from node to node; there are varying numbers of vehicles present at nodes; and there are varying-length (maybe zero) queues of pedestrians and vehicles waiting to access nodes.

Thus there are three types of latencies to manage in a mobility-on-demand system: pickup latencies, transit latencies, and dropoff latencies. The total latency for a trip is the sum of these. Users will care both about mean latencies and variances – since they want not only to minimize their trip times, but also to predict them with reasonable accuracy.

In general, the larger the number of vehicles in the network, and the larger the number of parking spaces, the shorter the queues and associated latencies will be. (You can always solve latency problems with capacity.) However, the costs, parking space demands, and road space demands will also rise. Therefore the operator's goal, in managing a mobility-on-demand system, is to meet user requirements for low-latency service without spending an unsustainable amount on vehicles and parking spaces. Software tools to facilitate achieving this goal are key elements of mobility-on-demand systems.

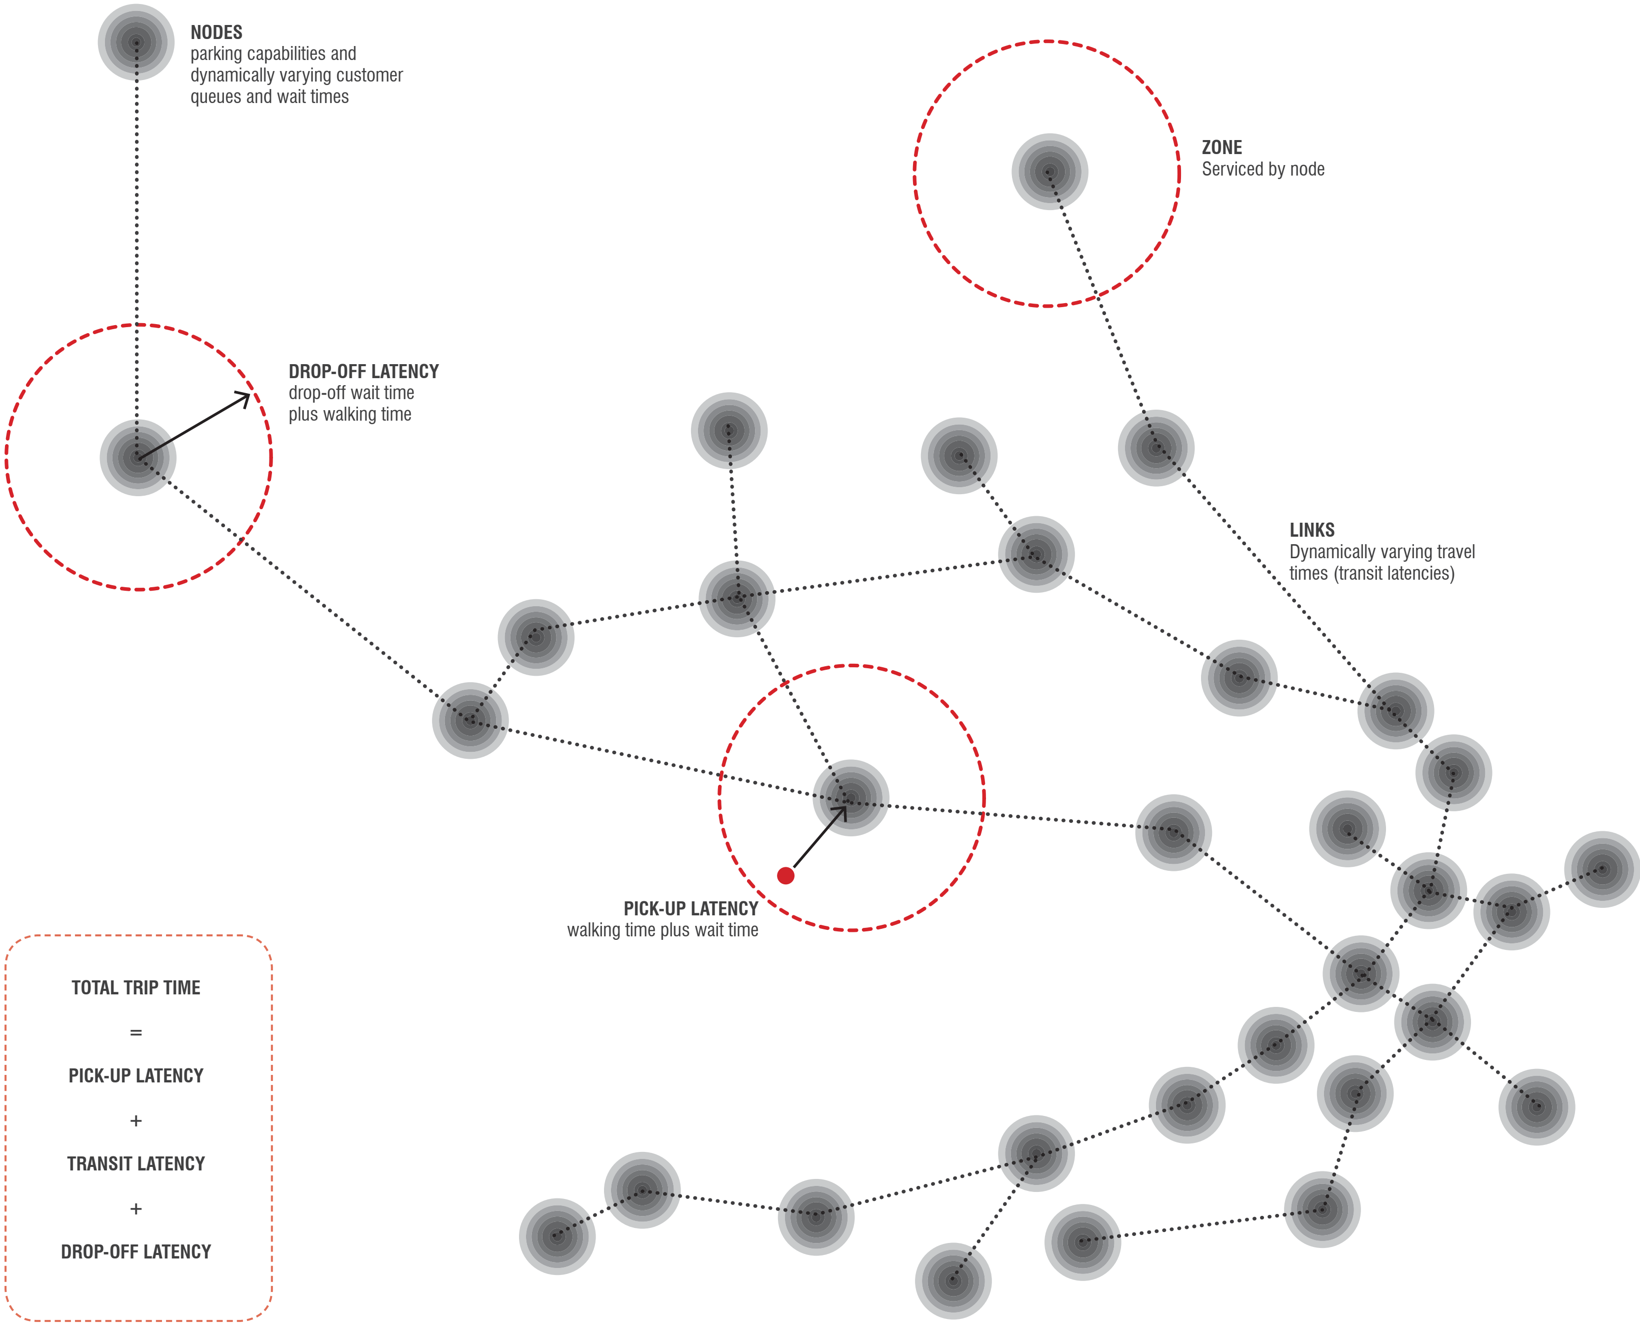
NODES
parking capabilities and dynamically varying customer queues and wait times
ZONE
Serviced by node
DROP-OFF LATENCY
drop-off wait time plus walking time
LINKS
Dynamically varying travel times (transit latencies)
PICK-UP LATENCY
walking time plus wait time
TOTAL TRIP TIME
=
PICK-UP LATENCY
+
TRANSIT LATENCY
+
DROP-OFF LATENCY

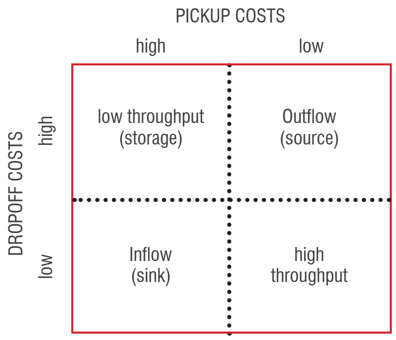

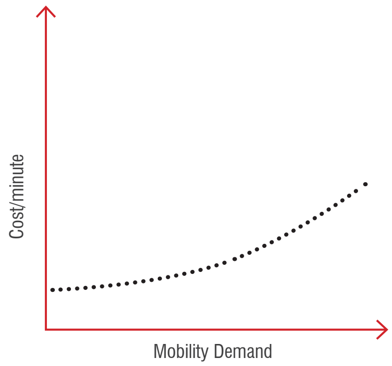

<< Managing vehicle distribution through node pricing

< Managing trip timing through peak pricing

> Troughs and peaks of bike availability experienced by the Paris bicycle system over 1 day

> Effect of dynamic pricing on desirability of pick up and dropoff nodes

Mobility demand manifests itself as queues forming at pickup points. A mobility-on-demand system must be designed to respond to this demand effectively (from the user's perspective) and economically (from the operator's perspective).

Cleverly managing the spatial distribution of vehicles in the network, as the spatial and temporal distribution of demand fluctuates, is the key to success. In other words, with a given stock of vehicles and parking spaces, the system operator must try to keep supply and demand of vehicles in optimal balance across the system.

Under certain ideal conditions – for example, in high-density, mixed-use urban areas with random distributions of trip demand – mobility-on-demand systems may be essentially self-organizing. In other words, the inflows of vehicles to nodes generally match outflows, so that there are never too many or too few vehicles at a location for the current demand.

In practice – as with other types of networks, such as electrical power, packet-switching, and delivery route networks – there will be spikes and irregularities in mobility demand patterns . This requires active management intervention – either by means of automatic control algorithms, by skilled operators who monitor and adjust the system, or some combination of the two – to keep supply and demand appropriately balanced. You don't want all of the vehicles on one side of the city when all the demand is on the other.

Since people make cost and convenience tradeoffs in their mobility behavior, and generally have some flexibility about when and where to go, much of this management can be accomplished through tools of dynamic pricing. If pickup price at a node is currently low, it will motivate users to go to that node, but if it is high, it will motivate users to seek a slightly less convenient alternative. If dropoff price is low it will attract vehicles to that node, but if it is high it will push vehicles out to alternatives. If price/time is low, it will encourage users to make their trip now, but if price/time is high, then it will encourage them to make the trip earlier or later .

If there is an available pool of appropriate labor, negative pricing may also be used. In other words, users can get cash or credit for moving vehicles to where they are urgently needed. This may be appealing, for example, to young people with time on their hands, the under-employed, and those who just want to explore the city or get some bicycle exercise.

Note, incidentally, that users at different nodes across the city may have different preferences for cost-latency combinations. In lower-income areas there may be a preference for lower-cost service with higher latencies. In higher-income areas, conversely, there may be a preference for lower latencies at a higher price. Pricing can also be used to implement public policy, for example by subsidizing the daily commute of low-income service workers to areas where they are needed .

However, pricing strategies may not always suffice to keep the system in an optimum state of balance between vehicle supply and demand. In this case, it becomes necessary for the operator to physically move empty vehicles from locations of current low demand to locations of current high demand. Obviously this is costly, and a management goal is to minimize it. Emerging techniques for efficiently moving driverless vehicles, such as virtual towing of trains of vehicles, and low-speed autonomous driving late at night, can assist with this. So (at least on a fairly small scale) can simply throwing vehicles on trucks.

It is not necessary to invent from scratch strategies and algorithms for balancing supply and demand in mobility-on-demand systems. The task is closely analogous to some well-known, extensively studied tasks such as airline fleet management, and delivery vehicle fleet management. There is a lot of existing theory, technology, and experience to draw upon.

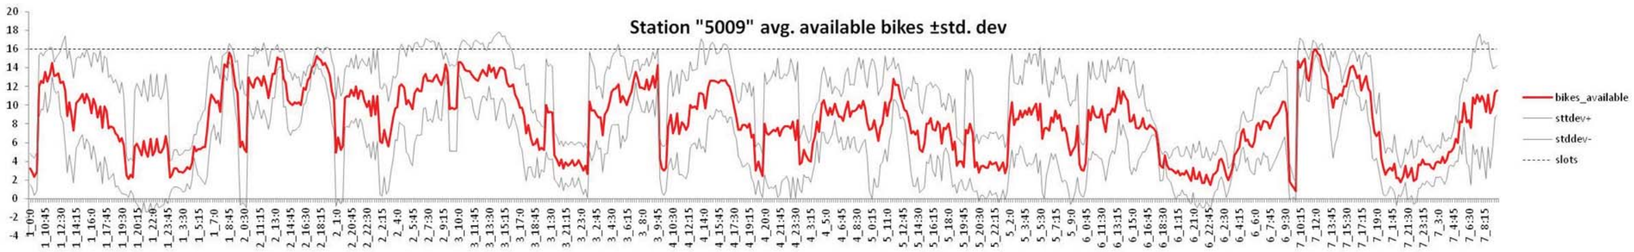

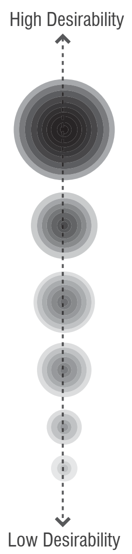

Node Desirability

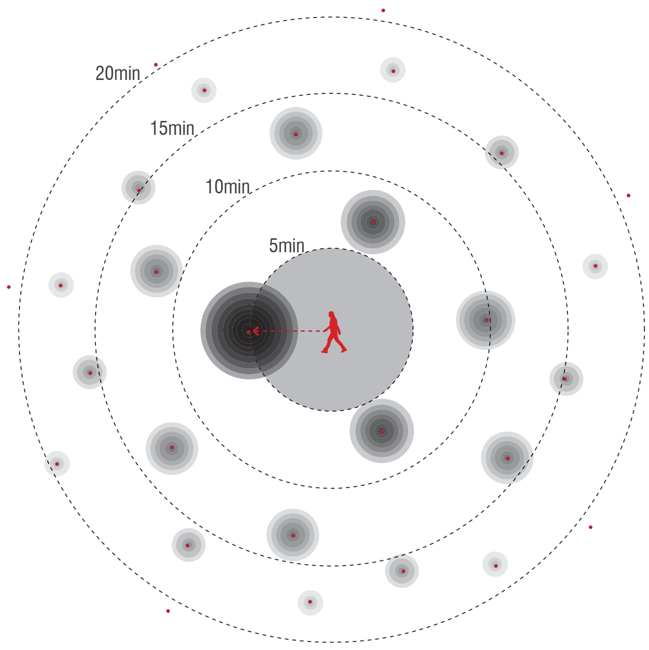

Desirability of nodes at equal prices

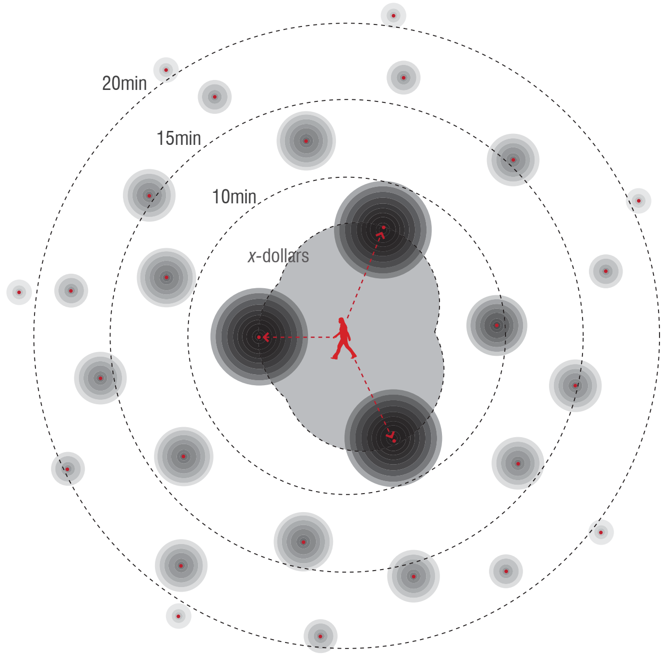

The effect of dynamic pricing on node desirability

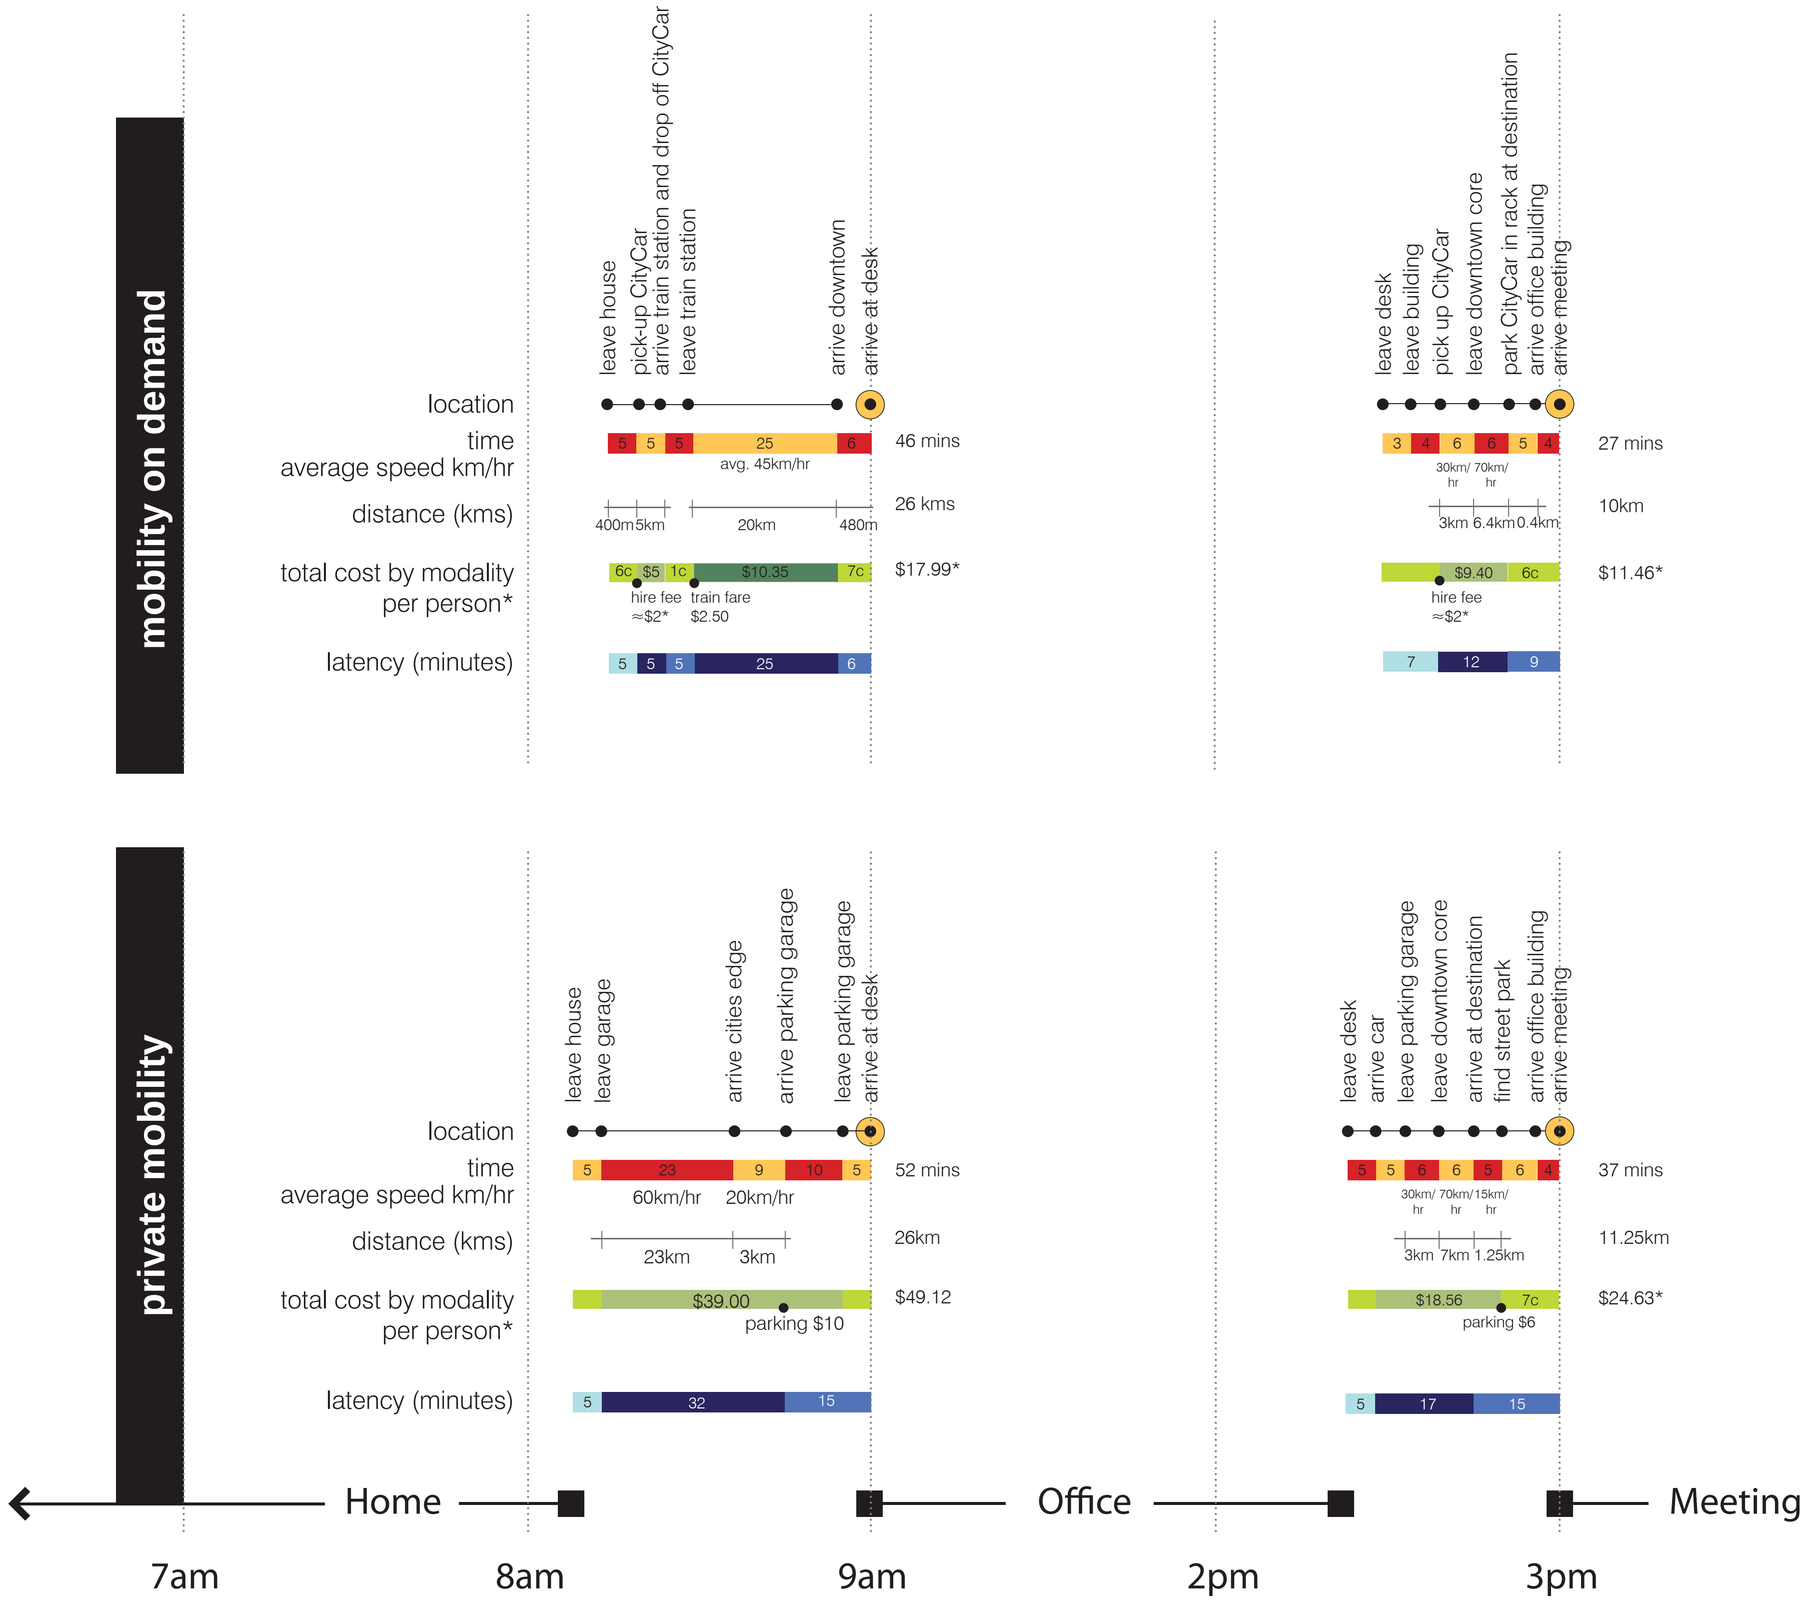
mobility on demand
private mobility
location
time
average speed km/hr
distance (kms)
total cost by modality
per person*
latency (minutes)
leave house
pick-up CityCar
arrive train station and drop off CityCar
leave train station
arrive downtown
arrive at desk
5 5 5 25 6
46 mins
avg. 45km/hr
400m 5km 20km 480m
26 kms
6c $5 1c $10.35 7c
hire fee ≈$2*
train fare $2.50
$17.99*
5 5 5 25 6
leave desk
leave building
pick up CityCar
leave downtown core
park CityCar in rack at destination
arrive office building
arrive meeting
3 4 6 6 5 4
27 mins
30km/hr 70km/hr
3km 6.4km 0.4km
10km
$9.40 6c
hire fee ≈$2*
$11.46*
7 12 9
leave house
leave garage
arrive cities edge
arrive parking garage
leave parking garage
arrive at desk
5 23 9 10 5
52 mins
60km/hr 20km/hr
23km 3km
26km
$39.00
parking $10
$49.12
5 32 15
leave desk
arrive car
leave parking garage
leave downtown core
arrive at destination
find street park
arrive office building
arrive meeting
5 5 6 6 5 6 4
37 mins
30km/hr 70km/hr 15km/hr
3km 7km 1.25km
11.25km
$18.56 7c
parking $6
$24.63*
5 17 15
Home
Office
Meeting
7am
8am
9am
2pm
3pm

Latency and cost comparisons: private mobility vs mobility-on-demand

* Cost Estimates in US dollars per vehicle kilometre

Victoria Transport Policy Institute Spreadsheet for Transport Cost Analysis, 2002

Total cost includes values for:

Vehicle Ownership	Congestion
Vehicle Operation	Road Facilities
Operating Subsidy	Land Value
Internal Crash	Traffic Services
External Crash	Transport Diversity
Internal Parking	Air Pollution
External Parking	Noise
Resource Externatilies	Barrier Effect
Land Use Impacts	Water Pollution
Waste	

1996 US dollars per mobility mode

Average Car	$1.65/km
CityCar	$1.00/km
Average motorcycle	$2.50/km
Roboscooter	$??
Bicycle	$0.42/km
Walking	$0.14/km

Total costs per day: private mobility vs mobility on demand

Costs	Private Mobility	Mobility on Demand
Time (mins)	134	113
Distance (km)	51.45	49.8
Cost (1996 US$)	$96.99	$44.20
Pick up latency (mins)	27	27
Travel latency (mins)	63	55
Drop off latency (mins)	44	31

< Latency, time and cost analysis between private mobility and mobility on demand systems.

Cost analysis based on the Victoria Transport Policy Institute spreadsheet for Transport Cost Analysis, 2002

All costs are in 1996 US dollars

> Callphone data provides 'census snapshots' of how people are distributed in a city. (Taken from senseable cities, Graz project)

>> The onboard location and guidance system on the roboscooter, provides the customer with data and directions to the nearest scooter rack, as well as monitoring the location of the scooter.

>>> Users of the mobility-on-demand system could employ their cellphones to check on waiting times or numbers of vehicles available

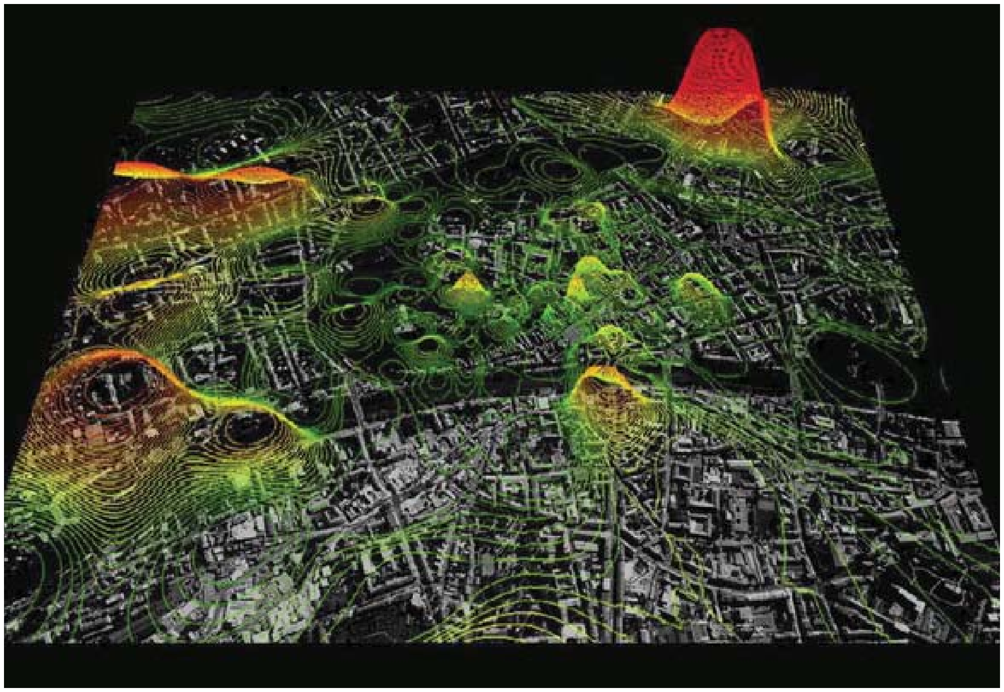

Obviously the effectiveness of mobility-on-demand systems depends upon having very good, spatially and temporally fine-grained information about varying patterns of mobility demand, and upon responding swiftly and appropriately to these patterns. (Traditional population density and trip origin/destination data, as used in transportation planning, may provide a useful starting point for approximately estimating demand, but it will not suffice.) There are several potential ways – maybe best used in combination – to obtain the necessary fine-grained, up-to-the-minute demand information.

The most obvious is to use the information generated by the mobility system itself. Nodes can keep precise track of queue lengths and actual vehicle pick-ups and dropoffs, and GPS navigation systems in vehicles can track vehicle locations. In this way, it is possible to build up very detailed historical pictures of expressed demand on the system, and over time, these provide an increasingly precise and reliable foundation for effectively responding to demand .

Cellphone operators now track the locations of handsets with increasing precision – either by their association to cell towers, or through GPS location. Aggregate cellphone location data provides realtime "census snapshots" of how people are distributed throughout the space of a city. This provides an effective basis for predicting emerging mobility demands .

Credit card transactions are recorded by time and location, so these can also provide a great deal of useful, real-time information about the distribution and activities of people in the city, and hence about likely patterns of mobility demand.

In general, mobility-on-demand systems draw upon large-scale, fine-grained historical databases to establish long-term patterns of mobility demand. They augment this with realtime sensor and transaction data reflecting short-term fluctuations and perturbations – due, for example, to special events or emergencies. They apply sophisticated analysis techniques to mine significant inferences from the data.

It should be fairly straightforward to establish the structures of the queuing and demand prediction models. These models will have many parameters, and initially the estimates of parameter values will probably be very rough. But, with experience over time, it should be possible to refine these values and thus develop very powerful and accurate models. Possession of these models will be a competitive advantage to experienced operators of mobility-on-demand systems, and lack of them will be a barrier to new entrants.

When a mobility-on-demand system is being planned, detailed databases of demand patterns, vehicle movements, and latencies do not exist. However, it is possible to simulate system operation in order to develop initial strategies for responding to demand, balancing the system, and minimizing latencies. Then, as the real system comes up, these strategies can be modified incrementally in the light of real data.

Mobility-on-demand vehicles are most effective when they are equipped with GPS navigation augmented with traffic density data. From the user's perspective, this enables efficient navigation to destinations. From the operator's perspective, it enables tracking of vehicle locations and provides realtime information about vehicle densities and speeds. Furthermore, destinations entered by users into navigation systems constitute "flight plans" that enable operators to predict parking space demands at arrival points and availabilities of vehicles, at these points, to meet near-future demands.

An even better option is to integrate mobility-on-demand systems with the emerging idea of personal mobility assistants (PMAs). PMAs are wirelessly networked, location-aware, handheld devices. They know about street networks, traffic conditions, transit routes, and transit schedules. They allow users, with minimum cognitive load, to plan and execute multimodal trips that may combine walking, mobility-on-demand, and transit – even when they are unfamiliar with the urban terrain that they are traversing.

The availability of high-quality trip planning facilitates sophisticated, dynamic pricing and the use of pricing to manage demand. Users can choose among combinations of vehicles, links, pickup and dropoff points, and overall latencies and prices. They may choose to optimize whatever combination of monetary cost, energy consumption, carbon footprint, and overall latency is important to them. In addition to minimizing resource use in this way, they might also want to maximize quantities like touristic interest or protection from the weather.

Location-awareness also opens up the potentially lucrative possibility of integrating location-based advertising, allowing users to plan shopping trips, combination with social networking, scheduling, and meeting co-ordination, and so on. Opportunities for innovative, add-on services such as these are likely to be important parts of mobility-on-demand business models.

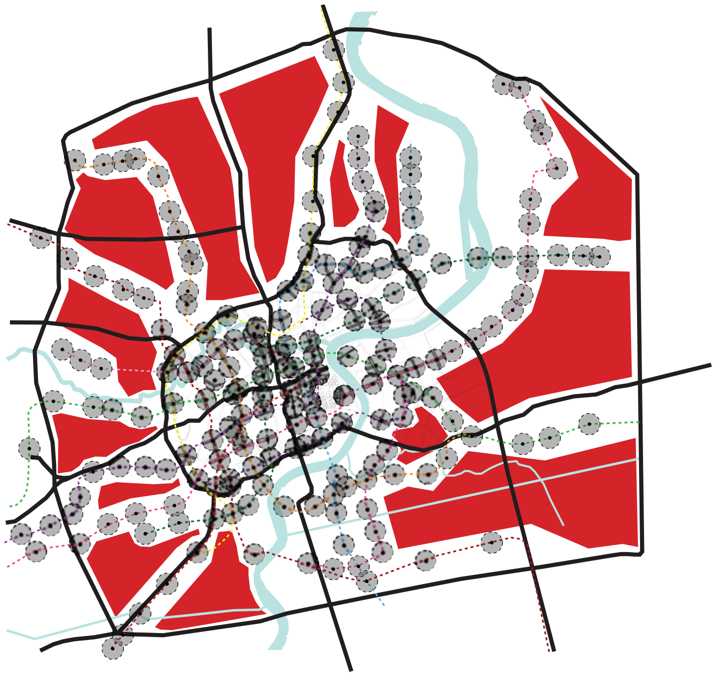

Mobility-on-demand systems generally are not replacements for transit systems. Instead, they operate effectively as partners of transit systems, and enhance the efficiency and attractiveness of these systems by solving the "first kilometer" and "last kilometer" problems.

In general, transit systems are very efficient for moving large numbers of passengers, at relatively high speed, between fixed points. Their difficulty is that boarding points are rarely exactly where you want to begin your journey, and dropoff points are rarely exactly where you want to end. You have to get to the embarkation point (the "first kilometer" problem) and from the dropoff point (the "last kilometer" problem).

Mobility-on-demand systems solve these problems by providing stacks and racks at transit stops.

One possible combination is with metropolitan transit networks, such as subway and bus rapid transit systems. Commuters might ride the suburban train home in the evening, pick up a vehicle at the stop, keep it overnight, and bring it back to the station in the morning. (There might be a price incentive to recharge the vehicle, at a home station, overnight.) At the city end, commuters might take vehicles from the station to the workplace and back again.

Another possibility is combination with inter-city high-speed rail or air transport, in which vehicles are picked up and dropped off at train stations and airports. This combines the long-distance speed and efficiency of transit systems with the short-range convenience of mobility-on-demand systems. And it eliminates the need to design mobility-on-demand vehicles to meet the requirements of long-distance, high-speed highway driving.

Transit systems are least efficient where population density is low and stops are sparse, and in off-peak times, when they have to move around large vehicles containing few passengers. In these contexts, through use of small vehicles and availability on demand, mobility-on-demand systems may cost-effectively substitute for transit .

Mobility-on-demand systems can also provide "virtual rings" to supplement radial suburban transit systems. In these systems, it is usually necessary to go into the center and out again to move circumferentially. The problem gets worse towards the periphery, as the radial lines spread apart. Mobility-on-demand stacks in suburban areas can enable efficient circumferential movement instead .

> Transport 'islands' in San Francisco. These areas do not have adequate transport facilities and are ideal places for mobility on demand systems as a way to augment the existing public transport network

< Transport 'islands' in the city of Shanghai


Low Access Transport Islands
5 mn walking distance radius
0m
200
1000
100
500
3000

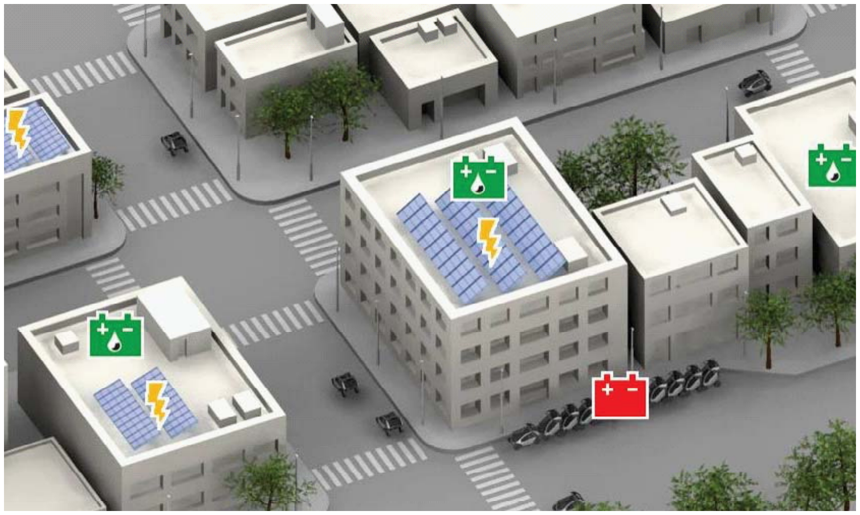

<< Flow of energy between solar panels on roofs to CityCar stacks

< The removable battery of the RoboScooter, makes it easy to ensure a fully charged battery at the start of each journey

> Carbon emissions and CO2 comparisons show environmental performance of vehicle types

The use of electric vehicles and bicycles eliminates tailpipe emissions, local pollution, and traffic noise. However, this does not necessarily reduce dependency upon non-renewable energy sources. This depends upon the source of electricity. If the source of electricity is old-fashioned coal-burning power plants, for example, then the shift to electric vehicles merely displaces (though maybe with at least some advantage) carbon emissions. But if the source is hydro, then carbon emissions are eliminated.

A general problem with today's electric grids is that they lack storage capacity. This makes it difficult for them to respond effectively to demand spikes, and it makes them unfriendly to clean, renewable, but intermittent sources such as solar, wind, and wave. However, since electric-powered mobility-on-demand vehicles are always connected to the grid when parked in stacks and racks, they throw a large amount of battery storage capacity into the grid. This opens up the possibility of vehicles buying and selling electricity – much as has been proposed for plug-in hybrids. Trading strategy would respond to current electricity prices and expectation that they would need electricity for travel in the near future .

For example, vehicles parked at homes could recharge inexpensively at off-peak times during the early hours of the morning, and might sell electricity back to the grid if they happened to be parked at home on a "sick day" during peak travel hours.

This also deals with the problem of intermittency in solar, wind, and wave generation systems. Cars can charge batteries while the sun shines or the wind blows, and then sell electricity when these sources are not producing. There is particular promise, in sunny climates, in combining vehicle battery charging with distributed, rooftop solar panels, since this minimizes transmission losses .

Charging and discharging batteries is not cost-free, and the costs of charging and discharging currently available batteries limit the immediate practical effectiveness of this attractive strategy. But improvements in battery technology will probably make it increasingly feasible.

Increasing political and economic pressure related to the geopolitics of energy supply, the need to reduce carbon emissions and global warming (to which gasoline-powered urban mobility is a major contributor), and the need to shift to clean, renewable energy systems, will create increasingly powerful incentives for local and national governments to support mobility-on-demand systems.

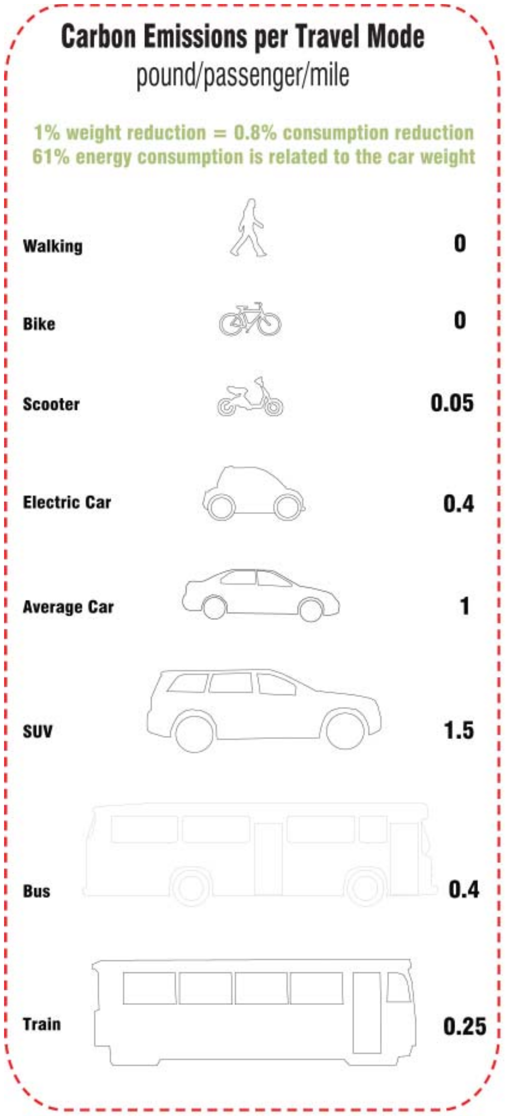
Carbon Emissions per Travel Mode
pound/passenger/mile
1% weight reduction = 0.8% consumption reduction
61% energy consumption is related to the car weight
Walking 0
Bike 0
Scooter 0.05
Electric Car 0.4
Average Car 1
SUV 1.5
Bus 0.4
Train 0.25

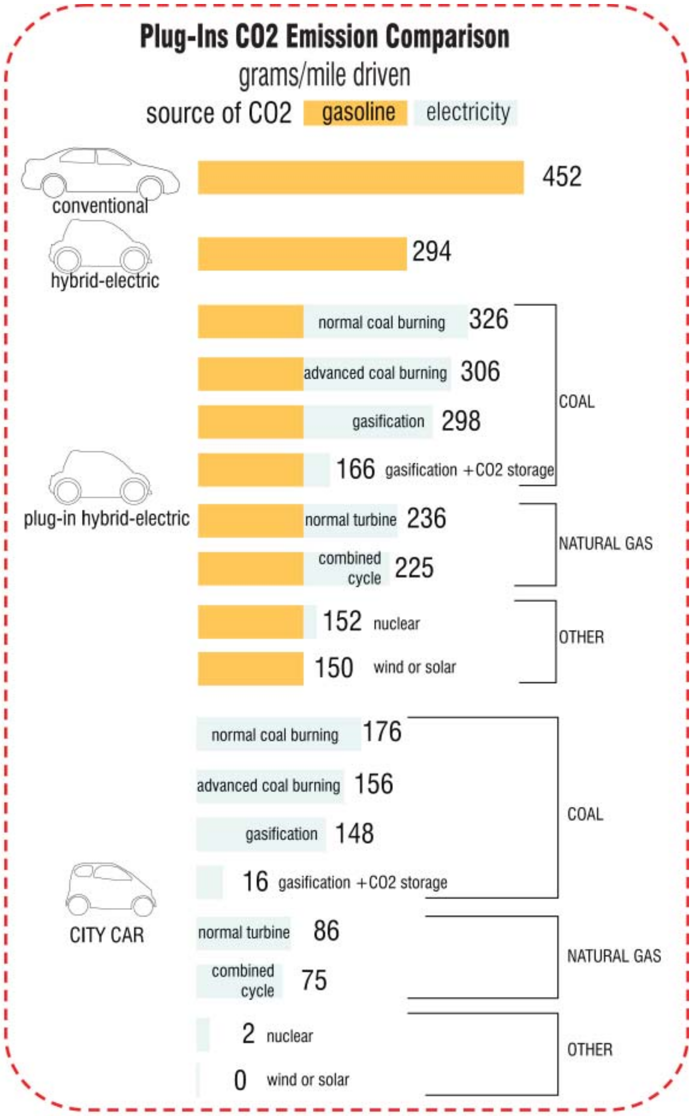
Plug-Ins CO2 Emission Comparison
grams/mile driven
source of CO2 gasoline electricity
conventional 452
hybrid-electric 294
plug-in hybrid-electric
normal coal burning 326
advanced coal burning 306
gasification 298
166 gasification + CO2 storage
COAL
normal turbine 236
combined cycle 225
NATURAL GAS
152 nuclear
150 wind or solar
OTHER
CITY CAR
normal coal burning 176
advanced coal burning 156
gasification 148
16 gasification + CO2 storage
COAL
normal turbine 86
combined cycle 75
NATURAL GAS
2 nuclear
0 wind or solar
OTHER

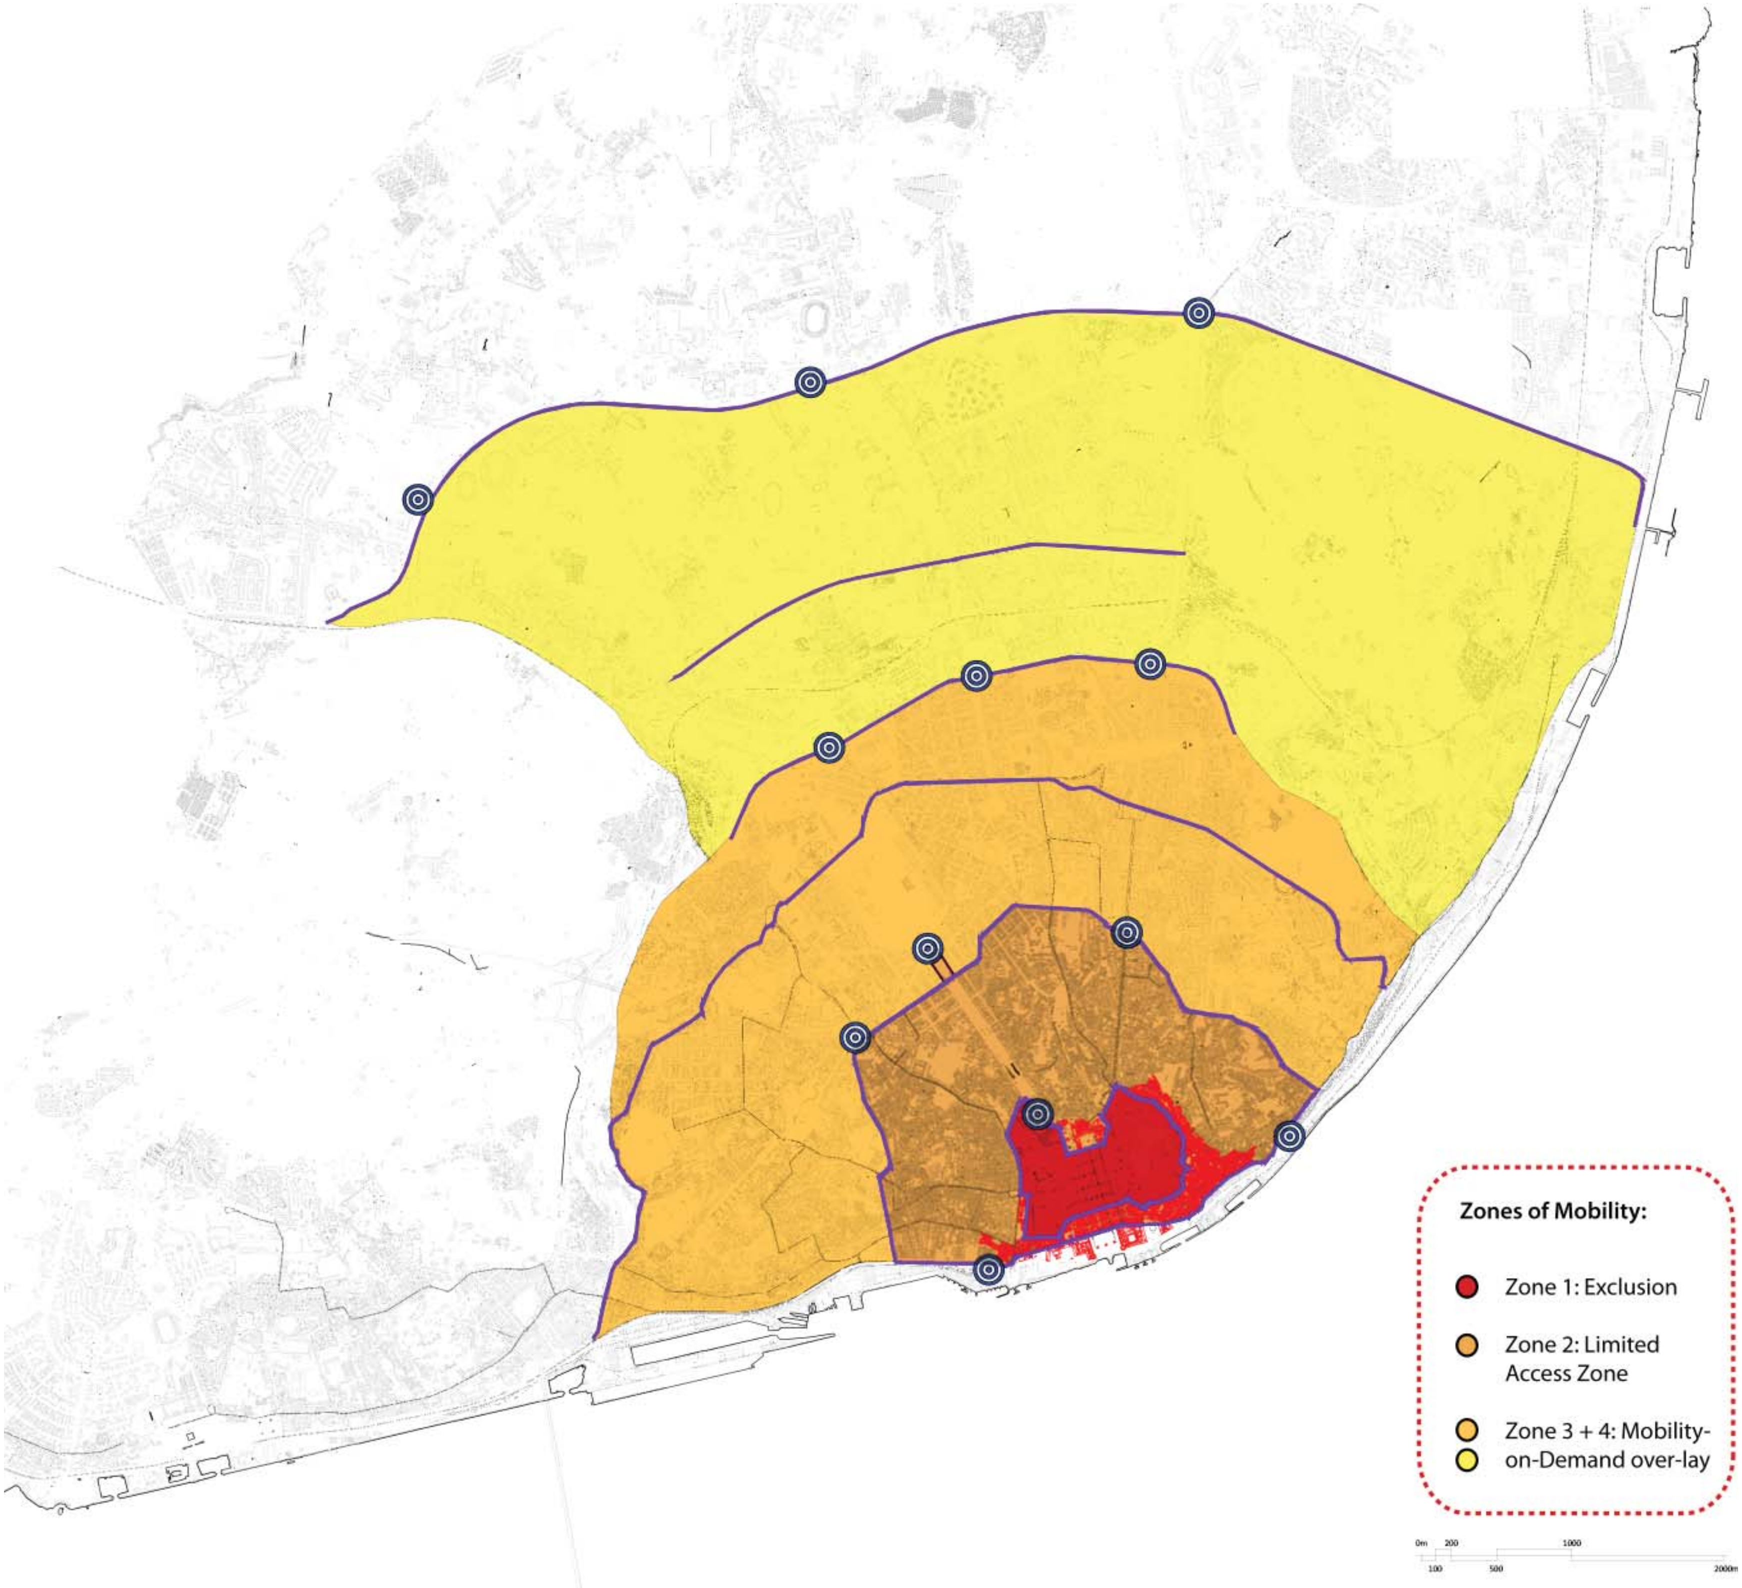
Zones of Mobility:
Zone 1: Exclusion
Zone 2: Limited Access Zone
Zone 3 + 4: Mobility-on-Demand over-lay
0m
200
1000
100
500
2000m

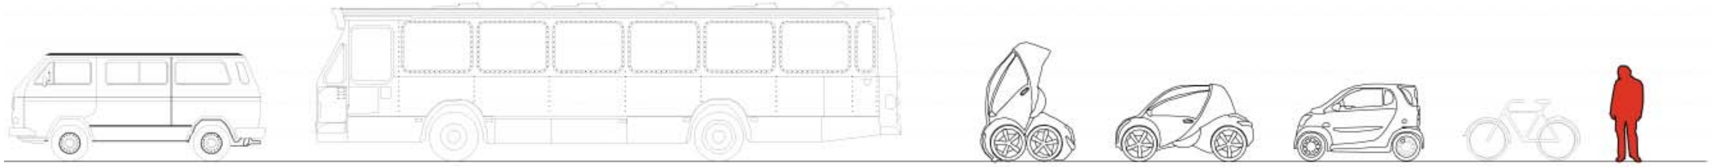

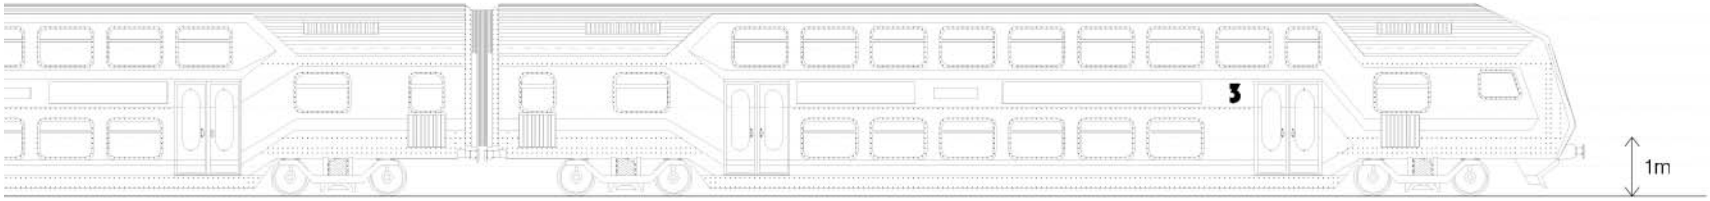

Currently, gasoline-powered automobiles weigh approximately 25 times the weight of the driver, and run at speeds of at least 150 km/hour. This combination of mass and velocity adds up to enormous inertia in crashes, and automobiles must be designed with safety systems to withstand this – greatly increasing weight, complexity, and cost, and reducing energy efficiency .

By making use of lower-speed, lower-mass vehicles, mobility-on-demand systems can enormously reduce overall levels of inertia in urban mobility systems – thus reducing energy consumption and embodied energy and the weight and complexity of safety systems, and potentially reducing overall road deaths and injuries.

Ideally, a mobility-on-demand system has a range of vehicles, from bicycles and Segways with low vehicle weight / passenger weight ratios to four-wheel, fully enclosed automobiles with higher ratios and safety cages, crumple zones, seatbelts, and airbags. Users can choose the combinations of costs, weights, and safety levels they want for particular conditions. Policy makers can also set general parameters .

As under today's conditions in most cities, there will be mixtures of vehicle sizes and weights on streets, and this will put smaller vehicles at a disadvantage in collisions with larger vehicles. But light vehicles forming mobility-on-demand systems will not need to mix with heavy vehicles under highway conditions, just as bicycles do not go on the freeways today. They will create demand and justification for higher levels of vehicle segregation by mass and speed, for example by establishing urban core zones that exclude or heavily penalize private automobiles and rely solely upon mobility-on-demand. And, by greatly increasing the percentage of light, relatively low-speed vehicles on the streets, they will reduce the average energy of collisions. The benefits to pedestrians (at 1/1 driver/vehicle weight ratio, unless they wear armor) are particularly significant.

< Zones of modality in Lisbon

^ Weight and size comparisons between transport modes clearly show the space benefits of a mobility on the demand CityCar

The private automobile has brought many benefits to city dwellers, but also many negative externalities. The effect of these externalities increases with automobile density. Streets become congested and slow to travel, noisy, polluted, and dangerous, and increasing proportions of valuable urban real estate must be devoted to car parking. From an urban design and quality of life perspective, the problem is to retain the benefits of the private automobile (particularly those that are realized at low automobile densities) while eliminating the negative externalities.

Mobility-on-demand systems accomplish this by providing equivalent or greater access to mobility with much more compact and benign vehicles, with far fewer parked vehicles occupying space, and with far fewer vehicles on the streets creating congestion.

The footprints of mobility-on-demand vehicles – even the electric cars – are much smaller than the footprints of gasoline-powered automobiles. Furthermore, they pack tightly together in stacks and racks. This allows parking compression ratios of between 3:1 and 8:1 for cars, and even more when two-wheelers replace cars.

One possible response to this gain is to pack many more vehicles into the same amount of parking space. This may occasionally be appropriate in areas where mobility demand is extremely high, but generally it will not be necessary. In most contexts, the introduction of mobility-on-demand systems, combined with the removal of traditional car parking, will result in freeing of on-street and off-street parking for other uses. Urban designers may take advantage of this to introduce greenery and other amenities into streets, to remove parking from piazzas and return them to public pedestrian use, and so on.

Although mobility-on-demand vehicles have lower top speeds than today's private automobiles, they can provide higher throughput in urban areas because they generate less congestion and provide through their navigation system for automatic routing around blockages and congested areas. Furthermore, since they are intelligent and wirelessly networked, they support sophisticated techniques of intelligent traffic flow management at merges, intersections, and constrictions. In general, they should be able to make highly optimized use of available street and road capacity.

> ^ Increased number of car parking spaces possible on a typical Lisbon block with the reduced footprint of the CityCar

> Possibility for increased urban amenity when car parking spaces can be reduced to 1.75m in width

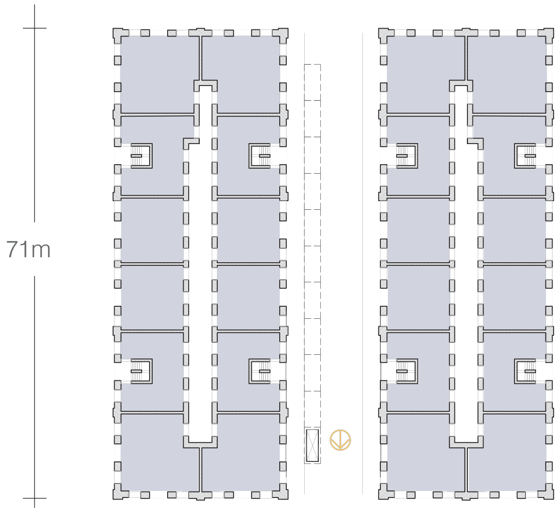

Typical block in Lisbon with 11
standard parking spaces

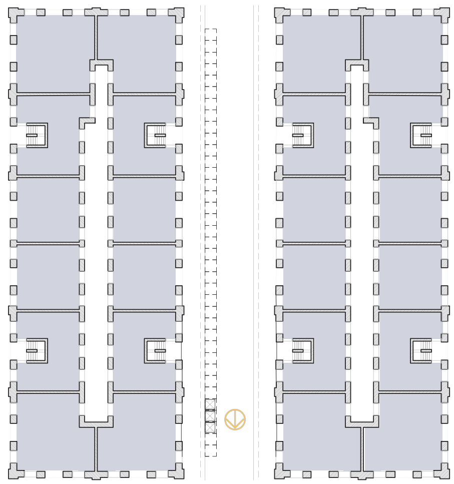

Typical block in Lisbon with potential
for 37 CityCar parking spaces

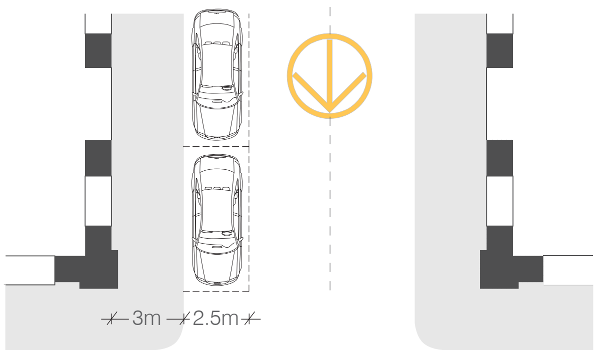

Typical sidewalk and parking
bay configuration in Lisbon

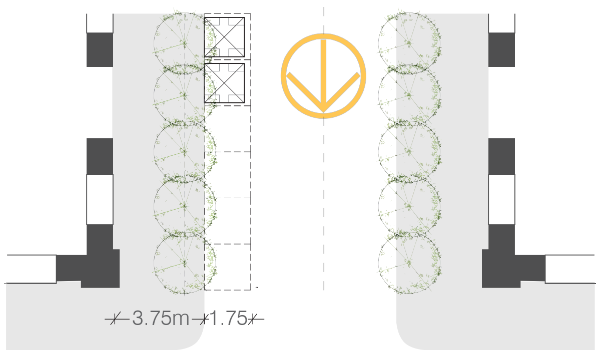

Potential for enlarged sidewalk
and street planting with reduced
CityCar parking bay of 1.75m