5 The Pulse of the City

> Tell me the size of a mammal and I can tell you, to about 85 per cent level, pretty much everything about its physiology and life history, such as how long it is going to live, how many offspring it will have, the length of its aorta, how long it will take to mature, what is the pulse rate in the ninth branch of its circuitry.
>
> —Geoffrey West, *Scale: The Universal Laws of Life and Death in Organisms, Cities and Companies*, 451

Let us see if Geoff West is right. The heart of a mouse pulses at 650 beats per minute and its average lifetime is some three years. Its life is thus just over 1 billion (1,025 million) heartbeats in length, and this number is the one that biologists have confirmed as a "law" that applies to most mammals. Applying this to an elephant whose heart pulses much more slowly, at 30 beats per second, then its average lifespan can be predicted as 1025 million divided by 30 beats, or 65 years. This is about right. Not bad, I hear you say, and West goes on to describe how this law and many other features of our physiology generalize across many species. In fact, humans mess up the "law" somewhat, since it only appears to apply to our average lifespans before the Industrial Revolution, when life expectancy was no more than 30 years.[1] The great transition has increased our lifespans dramatically, and as we implied in chapters 1 and 2, we stand at a threshold where medical advances might further accelerate this change. It would be nice to think that we could have a theory about cities with the power that West ascribes to our own physiology, and although much of that theory does resonate with how cities form and function, cities are much more complex systems than those in the animal world. As we will see in this chapter, however, there are deep and consistent regularities in the way cities function, both in the short and long term, but there are no fundamentally conservative properties such as those that we find in the biological world. The fact that in animal lifetimes, all hearts beat a similar number of times finds no equivalent when it comes to cities.

We have already demonstrated this, albeit indirectly, in the last chapter in showing that cities manifest allometry with respect to how the resources they generate or attract scale more than or less than proportionately as they grow in size. In fact, although the iron laws of animal physiology do not apply to manmade structures such as cities, allometry does provide a wonderful analogy that enables us to measure departures from such laws.[2] The economies of scale that generate increasing average incomes as cities get bigger reveal that certain processes intensify—and the one that is closest to our notion of the pulse is the speed at which we walk. That we seem to walk faster in bigger cities has been reported for many years. As far back as 1976, Marc and Helen Bornstein produced some fairly definitive evidence that we walk faster in big cities than in small, demonstrating that in places like Brooklyn people walked at speeds of 1.5 meters (nearly 5 feet) per second, twice as fast as in the small Greek villages they surveyed.[3] More recently, Walmsley and Lewis reported speeds of up to 1.68 meters (about 5.5 feet) per second for walking in London.[4] In 1970, Stanley Milgram (who we came across in chapter 3, where we introduced his work on small worlds) produced a more general summary of the pace of life in cities, characterizing how people undertook a range of activities at much faster speeds in bigger cities.[5] As a corollary, it seems vehicular traffic congestion also increases with city size. Using Google's drive speeds and routing algorithms, speeds per hour in US cities increase from an average of about 25 kmph (15.5 mph) in New York City, which has a population of around 8.5 million, to around 65 kmph (40 mph) for metro areas like Kansas City, whose population is some 480,000. If we look at speeds in the core of cities, anecdotal evidence suggests that in places like London's central business district (CBD), it is often faster to walk, since vehicular speeds hover around 10 kilometers (about 6 miles) or less per hour in the center. And in the biggest and oldest cities, these speeds seem to have remained pretty constant for several hundred years.[6]

Defining the pulse of the city is highly problematic. The word is used rather casually, and aggregating individual responses and taking averages for comparisons between different levels of activity is tricky. Much depends on cultural factors that cannot be quantified and on the historical evolution of the city with respect to its focus on activities such as entertainment and retailing. It is difficult to compare Las Vegas, for example, a city that is based on 24-hour gambling, with San Francisco, which reportedly has a population that on average sleeps less than any other in the United States. The activities during day and night are so different that comparisons are futile.[7]

However, before we explore the different pulses that define our cities and speculate on how the new digital world is yielding ever more data and information about how our cities function, we need to reflect a little on the way we look at cities with respect to how they change over time. Traditionally—and this is very clear from the last chapter—the way we have approached cities in the past is largely as though they are timeless; that is, frozen eternally at a moment in time. Although city planning does deal with the city of the future, the kind of master plans that have been the modus operandi of planners for at least a century assume an unspecified temporal future that will be reached at some point, but is invariably a convenient fiction to provide a focus. We do not yet understand, however, how the city forms gathered together as illustrations in chapter 4 actually evolved to the point where we could discern their structure, nor how the key processes of development that led to those forms will play out in the future. Past speculation has been solely confined to how cities have evolved over long historical time periods—over eras and epochs—in the manner used to explore long-term futures in chapter 2. Not that this was irrelevant, far from it, but it does not provide the basis for a requisite understanding of how cities evolve over decades and generations. In short, in most studies of urban form and function, there has been very little emphasis on urban dynamics.

Our thoughts about how cities function over much shorter periods, such as the 24-hour diurnal cycle, are not new—they go back at least to Chapin and Stewart's 1953 article *Population Densities Around the Clock*. Yet little could be done in terms of understanding these cycles until quite recently.[8] Richard Meier's attempts to construct a theory for urban structure based on thinking about social life in terms of communications theory provided a brave attempt long before its time.[9] But only when the digital world emerged, with all its sensors enabling us to track what is happening second by second, has the prospect of an entirely different way of thinking about cities emerged. Most of our focus in cities to date has been on densities that pertain to space, rather than intensities that relate to time. But this is changing rapidly, and in some respects this chapter, with its rather generic title pertaining to "pulse," redresses the balance—by thinking about cities in time, which is key to understanding their dynamics. In the past, our focus has been on cities in equilibrium, and although there have been attempts to examine abrupt change, such as catastrophes, chaos, and singularities that pertain to urban discontinuities, these have been over longer rather than shorter periods of time. Although complexity theory, alluded to in chapters 1 and 2, did change the emphasis on equilibrium to thinking about cities as being perpetually in disequilibrium in the short term and always being far from equilibrium in the long term, its focus has still been on change over decades or longer.[10] There have, of course, been attempts at looking at economic cycles pertaining to cities, which range from those associated with long waves such as Kondratieff cycles to various kinds of trade cycle, and we will address some of these in later chapters. But these have never been really linked to spatial change. The move to the very short term is a recent phenomenon, one that largely dates from the data gathered second by second at a very fine spatial scale, invariably using new generations of sensors and computers that have appeared only in the last 10 years. And thus the focus of this chapter is on exploring what is possible using these new kinds of data and dynamics.

Temporal Cycles, Flows, and Flux

When we examine how individuals and groups in cities function in time, we can classify their schedules and activities into nested cycles that define their behavior at different levels of aggregation. To an extent, the previous chapters have gradually focused in on more and more detailed kinds of behavior, starting from long waves and cycles characterizing technological change in our first two chapters, to more detailed implications of how cities change their form and function in the previous chapter. Economic cycles over generations and decades determine the state of the economy and the development of new technologies, and, within these, development processes determine how land use, transport, and economic activities change their locations as cities evolve over years and decades. Focusing our microscope on the finest temporal intervals, however, is where we reveal the detailed pulses, flows, and fluxes that characterize urban behavior, where the key cycles are diurnal—based largely on how our bodily functions relate to the presence of night and day—and determine the ways in which we work and play. Mobility is central to these processes, and in relating such movements to location, once again the notion of the network takes pride of place.

We can also characterize key cycles at slightly more aggregate levels—for example, seasonal change and its influence on behaviors. This makes a difference to the times and places where specific events take place, for many activities in cities are organized on a yearly basis, in particular education. In fact, although there are diurnal and seasonal rhythms particularly relating to work, entertainment, education, and so on, peppered within these cycles there are distinct one-off as well as repeating events that stand out as critical to the way cities function and are organized, especially in sport and entertainment. Many flows that can be recorded at very detailed temporal scales, such as second by second, or at least minute by minute, can almost be visualized as fluids, with their liquid flow being punctuated by discrete events. Thus, frequency of flow is key to an understanding of such changes, and in this chapter we will characterize a series of examples that define the contemporary city. However, our quest is to go well beyond this, and to imply how we might link all these together in generating a much more profound understanding of how cities evolve. Although our focus will be on the very short term, the short term becomes the long term if enough events and flows are recorded over ever-longer time periods. Thus, from the short term, we should be able to observe secular trends that manifest themselves over much longer time horizons. This is the challenge we face in linking small-scale change to long term, and short term functions to the evolution of urban form.

The flows that define our current understanding of how elements in cities relate to one another are almost entirely material and deal with movements, such as the journey to work or to other distinct functions such as retailing, commercial activities, education, health provision, leisure pursuits, and so on. Added to this, movements between economic activities, such as the transport of goods and commodities using many different networks, together with the flow of individuals to commercial activities, define a hierarchy of markets used in exchange. In turn, it is possible to see this constellation of movements as determining the demand and supply of goods, reflecting how prices are determined, and ultimately the way in which wealth is generated in an economy. As a consequence, some have recently invoked the idea of the network as the basis of the modern urban economy.[11] Movements take place diurnally, at discrete times, and seasonally. But layered on top of these patterns, land uses, individuals and households, businesses, as well as the many activities that comprise agencies and institutions in the public sector, all move and relocate for a multitude of reasons that pertain to the economic constraints or opportunities they have to confront. While these movements generate patterns of migration when they are aggregated to longer time scales, essentially each is composed of individual actions that take place at very distinct points in time, although some of these processes work themselves out over longer time periods. Traditionally, information central to the operation of the economy has been material in some form; prior to the Industrial Revolution, most information involved face-to-face contacts or letters that were passed using the technologies of the time—horse and carriage, dispatch rider on horseback, or complex forms of manual signaling, as used during the Napoleonic wars, for example.

Once the Industrial Revolution began, new technologies were quickly invented that enabled all the elements that comprise the city, as well as new innovations generated from these technologies, to change the basis of exchange and movement. The internal combustion engine, in the first Industrial Revolution, and the electric motor, in the second, continue to dominate the way we move around cities. But it was in terms of the flow of written information—data—that the changes we are currently in the midst of began. Even in classical times, information passed by messengers on horseback was the fastest way to communicate any nontrivial conversation over distances longer than a mile, while bonfires and smoke signals have been time-honored ways of passing "bits" of information since pre-history. Logical signaling systems such as semaphore were used in the early Industrial Revolution, and in the mid-19th century, the telegraph and Morse code were invented. Hard on their heels came the telephone, and together these technologies marked the first "death of distance" as new worlds were quite literally opened up due to this communication. Together with the railway, these technologies changed the world.

As we have pointed out in earlier chapters, cities would not have grown had such technologies not come into play; and, of course, this changed the basis of global trade. One of the reasons we consider that the real divide between the old world and the new began with the first Industrial Revolution is this development of both material and communications technologies, which generate their own positive feedback. The characterization of this era as the "Victorian Internet" by Standage[12] reflects that fact that our current digital age is simply an extension, a deepening, and an integration of changes that began more than two centuries ago, when the world first began to divide between previous eras of hardly any cities at all to a time when everybody will eventually live in cities.

Television was invented in the 1920s, but it was not until the development of the digital computer during World War II that the interactive manipulation of information really came onto the agenda. We have recounted a little of this history in our second chapter, and we will say more in later ones, but it was not until the convergence of computers and telecommunications, which reached its zenith in the 1990s, that truly interactive computation began. Moreover, computing at a distance became a reality with the invention of hypertext, and with its use on the Internet via the web, the impact of computers and communications on locations and interactions has become massive. There are now many forms of electronic interaction that take place between fixed locations where no material product is being exchanged. All of these are likely to have some impact on the way in which we behave with respect to economic and social interactions, but as they are largely invisible, at least at first glance, and because they are all conducted at the individual level without very obvious sources that aggregate this activity, it is hard to extract their meaning with respect to patterns of location and physical movement in cities. In short, these new ways of communicating data and information—indeed, as we will see, generating new kinds of information—must have an impact on how and where we locate and how we use the traditional functions of the city, which are translated ultimately into physical forms.

In particular, we need to define the range of such digital information. First and foremost, there is email, which is highly liquid. It is almost impossible to guess the number of emails sent worldwide each day—large quantities are spam or junk—but a conservative estimate (from February 2015) is around 200 billion per day. The number of Google searches per day is about 3.5 billion. We can repeat such superlative numbers for many media involving the Internet.[13] In fact, since so many of these media use the Internet and web browser technology to generate their content on web pages, the number of web pages accessed each day is not very meaningful. Breaking all this down to location, or rather to city-specific terms, is highly problematic. Even though many devices that are used to access this information are geo-enabled—their location can be fixed from geo-positioning satellites—there is considerable noise associated with this, since many users do not enable their location in any way. This is not likely to change in the near future, either, although most Internet providers do capture and archive such data for their own use. Much of this electronic data is from social media that intersects with the economy to an extent but is largely individual. When it comes to business-related data and interaction, the picture is even more confused—but in any case, the sheer volume of such data now is so great that there is no synoptic vision regarding how to make sense of it all, notwithstanding the fact that much of it is inaccessible in any case.

In the past, with rather small amounts of data on location and interaction in cities, we were able to aggregate the data quickly, easily producing flow tables, patterns of locational activity, and so on. Our understanding was limited by the data available, and thus theories of how cities evolved in terms of their form and function, usually at a specific cross-section in time, were manageable, while attempts were made to make such theory consistent with what was observed. No such theoretical superstructure exists for this new digital world. In fact, it may take a while for any to emerge, for the rate of change in what is happening in cities and society appears to be much faster than the rate at which we are able to absorb the new data and make sense of it via our traditional or even new ways of looking at cities.[14]

Notions about how we might represent space and time as individual profiles have been developed for cities before. The so-called space-time cube has been popular for many years. Credited to Hägerstrand's work[15] in the 1960s, with it individual travel profiles are visualized related to activities undertaken on daily basis. These profiles are richly descriptive, but there has not been much progress in stitching individual profiles together to generate more aggregate behavioral relationships pertaining to daily time budgets and travel patterns. The same is true for more microsimulations of city form and function. Ultimately, we have to contemplate the possibility of there never being any comprehensive aggregate theory of how cities form and function, given their current level of complexity. We may be doomed to a world where only partial theories are possible, where all our understanding is contingent on what we observe over short periods of time and at particular instants. In some senses, this is already the case. Notions about cities based, for example, on the "standard model" are now being partly eclipsed by new forms of behavior and new sets of values that have a very different expression in cities from the ideas we held in the past.

In his book *Radical Technologies*, Adam Greenfield begins with a vignette based on how many very different individuals access social media and use various sensor technologies in their daily lives.[16] This contrasts with a world at the turn of the millennium where no such behaviors were possible and we did not have access to the sorts of technology that are now widespread. The rate of change in this domain has been daunting, with the iPhone introduced just over 10 years ago (at the time of writing) in mid-2007. This device is in the vanguard of such change, with the latest estimate of smartphones in use numbered at more than 2.8 billion, about one for every three persons on the planet. There are now more mobile devices than people, levels of penetration that are still increasing very rapidly, with complete saturation—that is, everyone having such a device—possible within the next decade.

The implications for how we think about cities are very dramatic. Meier,[17] in some respects, anticipated this nearly 60 years ago in talking of cities as being environments where everyone "is continually bombarded by messages from other persons." But Greenfield's summary of what an individual does through their "normal" day catalogs this richness—a person wakes and answers email, checks website news, sends tweets, taps in and out of transit system locations, reads ebooks, accesses web information for work, downloads and uploads data at work, buys lunch with a credit or debit card, and so on and so forth. This rich diet of digital access supplies a locational and temporal record that could go some way to understanding how an individual behaves, which is part and parcel of the quantified self.[18] As we have implied, scaling this up to overall space-time patterns of demand and incorporating one-off events into this picture is quite problematic and has rarely been attempted. Add to this the patterns of city facilities supply provided by government and businesses, which are also controlled by digital technologies—often via individuals accessing them to enable management—and the problem of synthesizing all these options is writ large.

In the rest of this chapter, we will broach these questions, first with respect to the data, largely focused on transit and mobility, becoming available from routine sensors, and then explore the flows and networks that can be extracted from such data for cities of different sizes. We will broach, though hardly answer, the question of what social media can tell us about the future city, and we will finish with some speculation regarding what all this might mean for the city in this century and beyond.

Sensing the City and Real-Time Streaming

For the last 200 years, if not longer, it has been generally accepted that to alleviate the many problems that plague the city and to improve the quality of life of its citizens, physical change to the city's form (and hence, its function) has been the preferred and perhaps least intrusive way. To this end, the forms reviewed in the last chapter and changes to their physical planning, usually focused on arranging land use and transportation in different ways, have been the dominant approach. Moreover, we have assumed that for changes to urban form to have any impact, plans are implemented over long time periods—years and decades, rather than weeks or months—consistent with the fact that our observations of how cities function is over similar and rather long periods of time. We might consider such observations to be a type of sensing, although what data we have had available is collected periodically, with low frequency—every 10 years, for example, in the case of population censuses—and, of course, such data is usually gleaned from direct questionnaires administered manually.

This concept depends on separating the real city from any ideas, practice, and theory we have about it. It is the paradigm adopted by those who are guardians of those ideas, practice, and theory—who exercise control, management, and planning over the real city, while all the time generating data that is then converted to information so our understanding (hence, control) can be better executed. Ever since computers were developed, digital simulations of the city have been attempted, first as models of location and interaction at coarse spatial scales, and more recently as digital visualizations of cities in terms of their geometry in two and three dimensions; and, of course, data itself has become almost entirely digital along the way. What is different is that within the past 10, and certainly 20, years, digital devices—computers and sensors—have been directly introduced into the real city for purposes of planning and control to improve efficiency as well as generate better planning and quality of life. From this embedding into the physical fabric of the city have come new streams of data, which is often generated at very fine temporal intervals—second by second, or even finer. This data is a by-product of the digital mechanisms used to control the city, and in this sense it is a kind of "exhaust." It is often highly unstructured and thus requires special techniques to make sense of it, and it is usually voluminous—hence the term "big data."

This change in the paradigm involves a switch in focus to much shorter time intervals associated with how cities work. Data gathered automatically using such sensors can be used in some cases for real-time management, as in the variety of control centers that have mushroomed in the last decade in the biggest cities, where traffic, in particular, as well as emergency services need constant monitoring and intervention so that such systems work smoothly. This new paradigm is often referred to as the "smart city," but other terms have been used: the information city, virtual city, digital city, electric city, and, earlier, even the wired city. In essence, this implies a massive change in focus from long term to short term, from large scale in spatial terms to much finer scale in individual terms, and from strategic to routine in terms of management and control (although, as noted, the short term eventually turns into the long term if enough data is collected over a long enough period of time). This represents a very different use of computers from that traditionally used to model cities and transportation to inform better planning. The fact that we can use the same devices to understand the city as those that generate data from it and control its functions is yet another demonstration of Turing's insight that the computer is "the" universal machine.

We can divide the way we sense the city into three regimes. First, we can introduce sensors to monitor the performance of built and natural environments we might condition in some way. This is physical in focus, for many of these sensors are organized to pick up the operations of natural and manmade systems and do not involve any human intervention other than periodic maintenance to capture the data generated and to ensure they are working. Such systems range from inductive-loop traffic detectors, which have existed in analog form for many years, to new forms of energy monitors that are part of the quest to automate buildings and the home. In contrast, there are also social sensors, which are devices we use ourselves. The entire panoply of social media is configured around our devices, which enable us to get access to this virtual world. In short, we decide when they are used on an event-by-event basis. Although the line between the social and the physical can be blurred, if we program our machines to perform routine tasks even at the individual level, the distinction is obvious. Somewhere in the middle are sensors we activate in a routine fashion. When individuals are scaled up to very large aggregate populations, such as those accessing a stock market, using a transit system, or even watching television, then such data is easily seen in real-time streams, despite the fact that it is made up of many individual objects and people. Another way of thinking of such data is as being passively or actively collected; the differences between these two methods can be profound with respect to our ability to use such data in making generalizations.

Passively generated data pertaining to manmade and natural objects usually generates information about different spatial locations. Movements of automated vehicles can be tracked, but it is data actively generated by ourselves that is more likely to be associated with network flows, in particular ways in which our devices are able to track our movements. In fact, we will not focus in this chapter on every possible kind of social and physical media developed over the last decade that yield information about where and when activities in cities take place and the particular individuals involved. We will examine short text messaging as an example of social media and movement data from smart card activations, which represent a blend of the social and physical. We will also look at the automated generation of data about vehicles, such as subway trains, where the vehicles themselves or their operators, not their users, activate the system in event-by-event mode.

Since so much data is now being generated routinely as real-time streams, and the variety of such data is quite wide and somewhat idiosyncratic, various portals have emerged that attempt to pull very diverse data together in the form of dashboards.[19] These are popular with city authorities who want an instant and immediate picture of what is happening with respect to more than one data stream, and they are easy to build at a crude level, for there are now many interfaces to such real-time streams or to archives of such data that are made available with short time lags. In figure 5.1(a), we show a very simple dashboard that pulls together real-time streamed data for London, and in figure 5.1(b) the web-based London Panopticon, which provides camera images of what is happening at the eight compass points around high-profile landmarks. All of these, of course, are in the public domain and can be accessed easily, but a dashboard collects and displays them in relatively comprehensible form. If one watches such a dashboard or portal for more than a few moments, trends can be observed. But the ones shown in figure 5.1 have no intelligence, although more useful analytics are being rapidly developed. What these do at present is provide some sense of how such data can be mobilized and how it is being used for short-term, often crisis-related management. As such, these dashboards and portals provide an intuitive synthesis of the "state" of a city, in comparison with those examining much more abstract data on less frequent cycles.[20]

There are now many online systems that enable us to extract data from cities in real time. Even map interfaces such as Google Maps and Google Earth allow data to be captured at various levels, which can then be used to explore a)

b)

Figure 5.1

Organizing real-time streamed data: (a) in a city dashboard (http://www.citydashboard.org); (b) in a web-camera portal (http://vis.oobrien.com/panopticon).

cities. NASA night lights data from satellite sensors has become popular, and although this is not quite available in real time, it will be soon and will then allow all kinds of real-time applications. Virtually everything that can be sensed and digitized in cities is now being processed, and the rate of change in addition to such real-time streaming is now so great that this section of our exploration of future cities is inevitably the most incomplete in this book. The data deluge is so massive that it threatens to overwhelm us. Everything about the digital age is dominated by superlatives, with that associated with data being the most dramatic. Current estimates[21] (from December 2016) suggest "90 percent of the data in the world today has been created in the last two years alone"—at some 2.5 quintillion bytes of data a day! The challenge of dealing with such big data, even at the city level, is enormous—plagued by problems of confidentiality, lack of structure in the raw data, missing data due to noise in sensors, diverse and often inconsistent attribute tracking of an individual's social media, and a host of issues pertaining to the general invisibility of important data such as email, web access, and so on. We will focus on all these issues in the next section, where we explore typical examples of the diurnal cycles that dominate the structure of city data.

Rhythms of Transportation: Demand and Supply from Online Data

The basic rhythm of the city is the regular and routine movement of people who work for a living, usually in cities in the developed world, five days a week. People mainly start work in the early morning and finish in the early evening, notwithstanding that there are deep cultural differences, reinforced by levels of prosperity and culture associated with differing social contexts, that in turn affect these patterns. For each working day (usually a weekday), there is a short, more intense peak of travel in the early morning, a slight decrease in the trough that occurs in the middle of the day, and then a more elongated peak in the late afternoon to early evening. In big cities, there is often a late-night peak, which is much smaller but nevertheless quite distinct, associated with entertainment venues. Weekend days are quite different in that there are no morning or evening peaks, but usually an increase toward the middle of the day and a long tail falling off as the evening approaches. There is sometimes a peak on Saturday nights. Seasonal changes confound this pattern somewhat, particularly with respect to education when school holidays start and end. The months in the middle of the year are mostly associated with holidays—usually July and August in the northern hemisphere—and these affect the volumes, while severe (hot or cold) climates can also alter the diurnal profile. One-off events show up in various profiles, particularly related to sport, concerts, and festivals, and these can be factored into different layers that separate routine from less regular patterns. In fact, there are many different levels of routine events. In some sense, the peaks themselves are one-off events based on individual travel decisions but are intensely clustered to become routine, while regular entertainment events are often organized on a less frequent but continual basis; and then there are simply events that are not predictable but have to be factored into the picture.

If you look at transit systems in the biggest cities, most of these have now developed smart card technologies that enable automated payment. This is largely because such systems tend to move very large numbers of passengers, and automation clearly speeds the process. Sensors enable these volumes as well as their related payments to be easily captured, sometimes for real-time control, but always in archived form so that policy analysts can gain a better understanding of the system in hindsight and develop appropriate policies for short- and medium-term control. These kinds of system are usually based on registering the start of a trip by a tap-in on a sensor and the end by a tap-out, but there are various more sophisticated systems that enable users to travel on more than one mode for the same price if within a fixed maximum time. When passenger data is captured in this way, it reveals the demand for transit as compared to the supply (which is the positioning and direction of travel of the vehicles themselves). Figure 5.2(a) shows the aggregate tap-in and tap-outs over the whole working day for the use of tube trains (subways) in London on a typical November day in 2010. The two peaks are clear, and the distinction between tap-ins (entries) and tap-outs (exits) is lagged, with the differences between the profiles showing the rise and fall of each peak. In figure 5.2(b), we show the diurnal profile for every one of 270 tube stations, which reveals that the routine features of the aggregate travel are clearly mirrored at more local levels, indicating very substantial order in how the network performs and the flows that characterize movement on it.[22]

This kind of data provides a clear picture of the demand for transit, but there is now routine data available online and in real time for the supply of transit—the positioning of the trains, how close they are to their predetermined timetable, any problems that pertain to their operation, and so

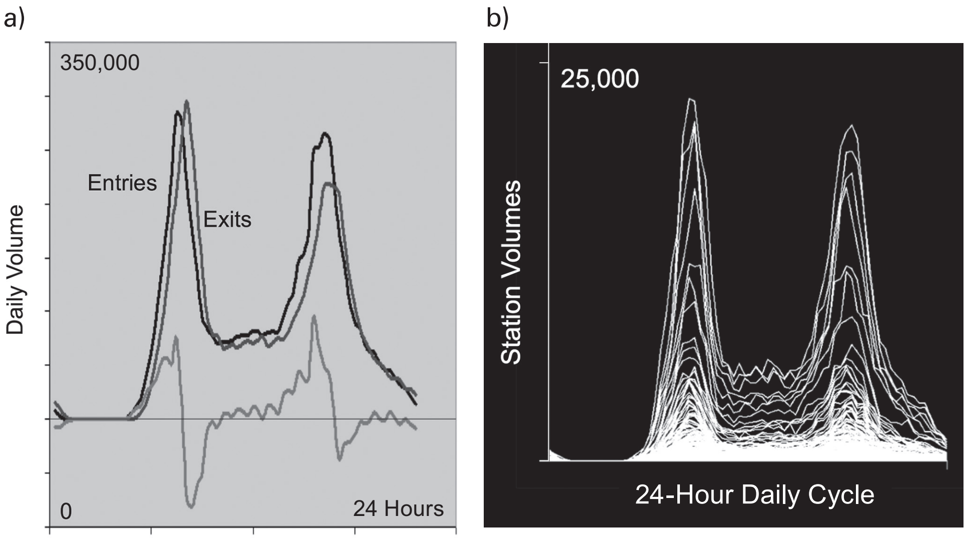

Figure 5.2
Diurnal passenger volumes on the London Tube: (a) aggregate diurnal flows; (b) diurnal flows at each underground station.

on. Much of this information is used in communicating with passengers if trains are delayed or disrupted. The supply data is entirely automatic, collected through the operation of the trains. This is what we referred to earlier as physically structured automated data, not dependent on any human input for activation or transmission, unlike the smart card passenger data, which is individually user generated. In fact, these vehicle sensors are by no means as reliable as those used for passenger demand inputs, and they have to be reset frequently to ensure continued operation. In short, there is more noise in this data. Data on the position of trains can be categorized by the tube lines, and in figure 5.3, we show a snap shot of four months' worth of weekly data colored by each tube line and accumulated for each time during the day of the week.[23] Much the same distribution of peaks and troughs during the working week that we see in the passenger demand are reflected in this data. But from these pictures, it is quite hard to distinguish any obvious seasonal change, other than the weeks when there are school holidays, where the total numbers of trains are clearly reduced, presumably according to predictable lower demand. In fact, the supply position in terms of trains has much less flexibility than that in the passenger demand data, simply due to the fact that individuals determine demand, while trains of different

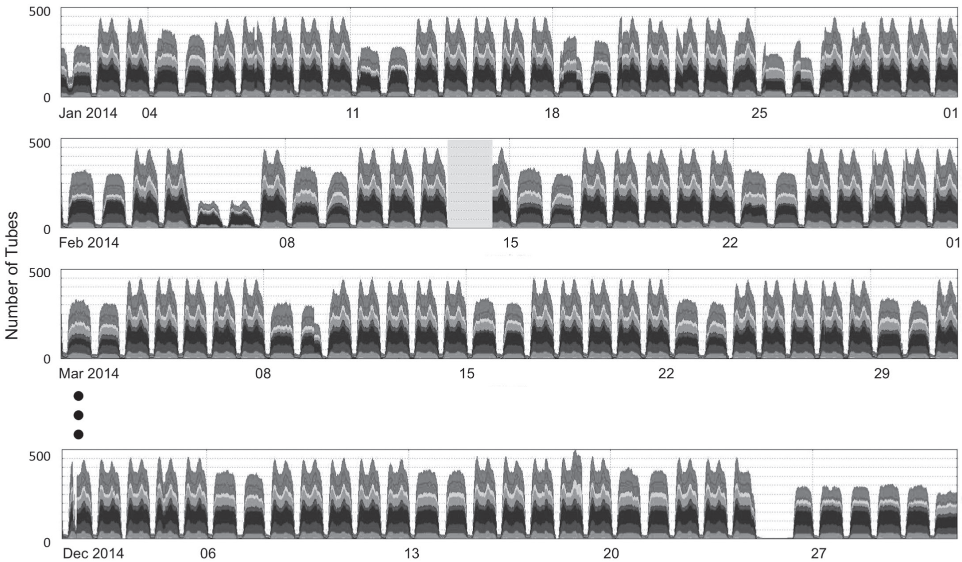

Figure 5.3
Train volumes on the London Tube over different months in a year.

sizes cannot be subdivided at will, hence there is much less flexibility in their deployment.

These patterns pertain to every system or sector that varies its functions in time and space. Thus all our illustrations here are applicable to other data streamed in real time and associated with the actions of individuals, ranging from retailing to social media. For transit data, we can examine how locations change over time—the hubs or stations where travelers connect up with trains—and how the flows between locations vary during the day. In figure 5.4(a), we provide a snapshot of tap-ins and tap-outs in the morning peak in central London, while in figures 5.4(b) to 5.4(d) we show the flows between stations at the morning peak, the midday trough, and the evening peak. Cities, as we know, are fluid systems that are in continual motion, and the best pictures we can use to communicate their dynamics are animations of how locations and flows change. We cannot use the printed page to illustrate changes over 24 hours—instead, we take three time instants in figure 5.4(b) to 5.4(d)—or for longer time periods, but a movie, created by Jon Reades from digital passenger demand data, is available online.[24]

Some rather good visualizations also exist for other cities, in particular for Lisbon, reminding us again that cities are like organisms, with traffic

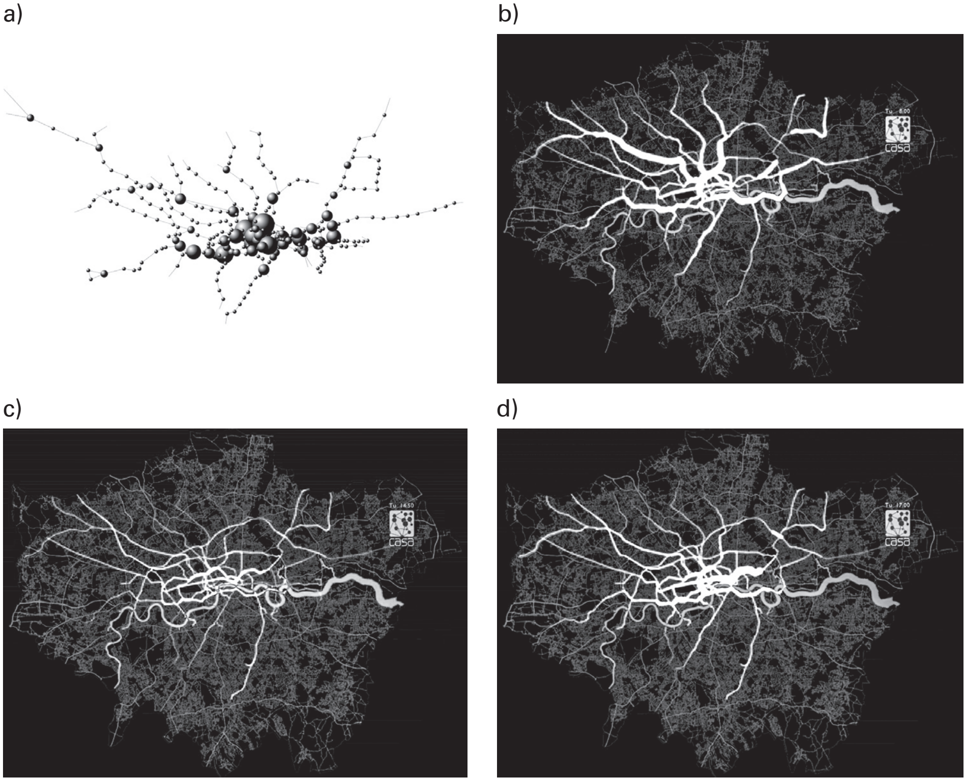

Figure 5.4
Locations and flows during the daily cycle: (a) aggregate flows at station locations; (b) flows in the morning peak, 8am; (c) flows in the noon hour; (d) flows in the evening peak, 6pm.

acting as their life force, these flows being similar to the way blood flows through the human body[25]; indeed, every flow of this kind one can imagine provides a similar demonstration of the many processes that encapsulate how we interact with one another to enable the city we live in to function in a workable, perhaps even optimal way. When the city becomes unworkable, congestion takes place, and the flows in these animations clog—the analogy to the human body is one of blood clots forming, traffic jamming. The implications for thinking about cities in these terms are fairly obvious: we need to design cities that do not clog but gain the advantages of traffic that flows fast and efficiently. This is a balance that is forever changing. It involves a complex spectrum of patterns whose variety and volatility imply

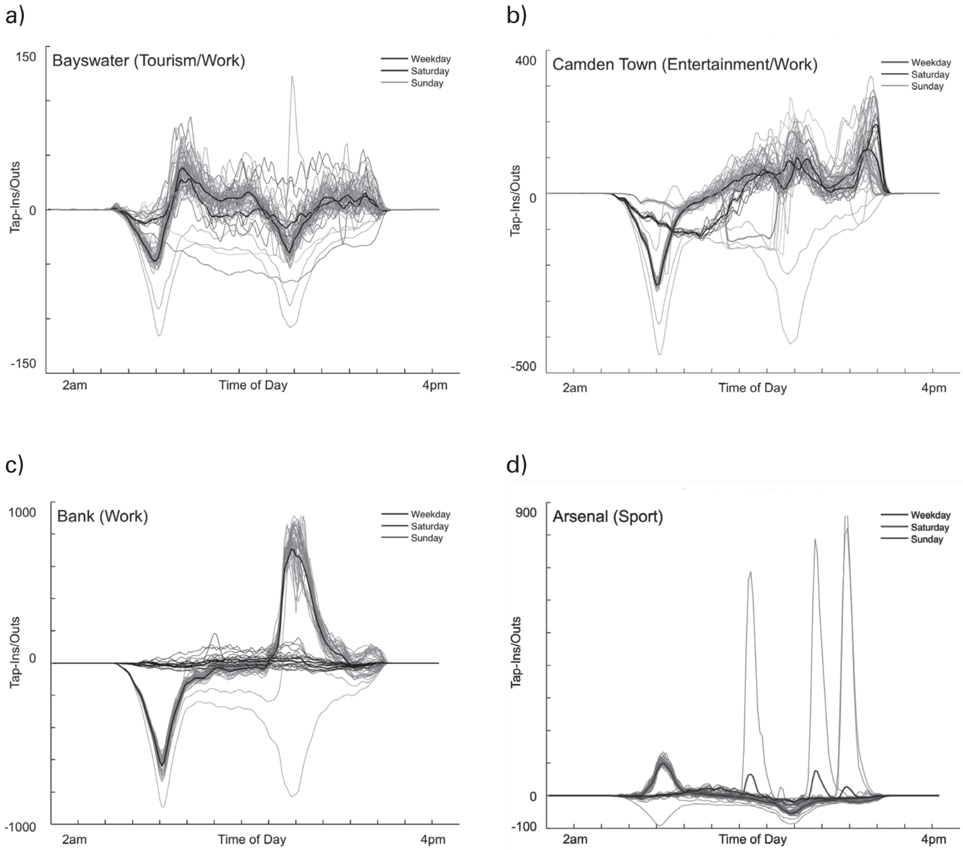

Figure 5.5

Temporal profile–based tap-outs at typical London subway stations: (a) Bayswater; (b) Camden; (c) Bank; (d) Arsenal.

a healthy functioning system analogous to the rhythms that define bodily functions.

When we examine these patterns in terms of locations, which in this case are the volumes associated with the subway stations, we produce a very detailed picture of the remarkable heterogeneity of travel behaviors associated with different places. In figures 5.5(a) to (d), we show four subway stations in London in terms of their weekly and weekend diurnal profiles, and it is clear these are very different. These are all shown as deviations from their baseline flow in terms of tap-outs, so they cannot be compared in absolute terms, but what is clear is that each has some semblance of morning and evening peaks. If you know these areas, then entertainment dominates (a) and (b), work (c), and sport (d), the last accounting for the three spikes in this graph relating to football games that happen at the Arsenal.[26]

Joining these patterns together to provide a synthetic view of the city will give us a better understanding, but there are many pitfalls along the way. For example, we have good data on the demand for travel by passengers and the supply of trains that will meet this demand in the examples we have introduced. But to really make sense of the way the transit system operates, we need to examine how passengers and trains are connected. If we are interested in disruptions and failures of equipment on the system, then we need to figure out how to connect passengers to trains. Subway stations can be very convoluted in terms of their layout, and there are many ways of getting from a tap-in point to a train; passengers have different cognitive abilities to make such connections, but if we seek a detailed and accurate picture of how a passenger is disrupted on a particular train, we need to be able to observe this link. In fact, in the London tube, there are very specific restrictions on tracking people using Wi-Fi, which is not allowed at present without special safeguards because of privacy concerns. Add to this the fact that the system is by no means foolproof with respect to tracking demand, because often barriers or gates are left open, and there are several groups of passengers who are not penalized if they do not tap-in or tap-out, making the problem of getting accurate data on disruption quite serious. It is not our intention in this book to inquire into the detail of big data of this variety, for we only wish to emphasize just how much this is changing our perceptions of what we consider important in cities and, in particular, in future cities. Yet although this revolution has the potential to provide us with a new understanding of location and interaction in cities and is giving us a massive variety of individual data pertaining to behavior, it raises many new problems that need to be resolved.

Cities as Networks of Social Media

To provide a comprehensive picture of the various pulses that characterize the rhythms present in the city at different scales and over different time periods, we probably need to consider some sort of classification of the city into different activities. This would let us integrate the various signatures of these activities so that we could see how they coordinate or conflict with one another. No such discussion, however, is yet possible, and all we can do is produce fragments of the various profiles that describe how the city works in the short term as well as implying that such profiles extend to the long term. In the last section, most of the activity associated with transit pertains to working, schooling, and possibly shopping, but all the activities that define a city are tied up in these profiles, as figure 5.5 makes very clear. There is now substantial information available pertaining to retail transactions, but most of this is proprietary and its spatial and temporal distribution is hard to visualize because there are few attempts so far to integrate such data. Where it is mainly used is in marketing and determining the location of new stores that maximize the profits of operators.

However, we are able to show data pertaining to credit card transactions in the city of Barcelona for a whole year (2011). Carlo Ratti and his group have produced a fascinating visualization of credit card flows for Spain from this data, which is available as a YouTube movie.[27] The annual cycle does not have the same sort of peaks that the transit cycles in the previous section show, but what is possible are visualizations of the patterns of demand and supply for various zones into which Barcelona can be divided. We show the flows for the maps of purchases where individuals reside and expenditure and revenues where sales take place in figures 5.6(a) and (b), respectively, with the mismatch between them relating to the differentials in spatial purchasing associated with the use of these cards. However, this does not map into the actual flows. We show samples of the aggregated expenditure and revenue flows for the top 30 locations in Barcelona in the abstracted circular graph in figure 5.6(c), where we plot the top 2 percent of flows in the darker gray colors. As you might imagine, there are flows from every place to every other place, and it is a challenge to visualize them so that we can extract informative patterns in the data.[28] Thus our illustrations are highly impressionistic, but it is not our purpose to conduct a detailed analysis here; we simply give the reader some sense of the kinds of exploration that are possible. We also have a map of the density of tweets over a week for the same city, produced by Neuhaus,[29] which we show in figure 5.6(d). We will pick this up again when we begin to explore social media in a bit more depth below.

Social media has become the collective term for online interactions that are generated by individuals communicating information about themselves or about the world at large from their own devices, often mobile devices such as smartphones. Such communications involve everything from

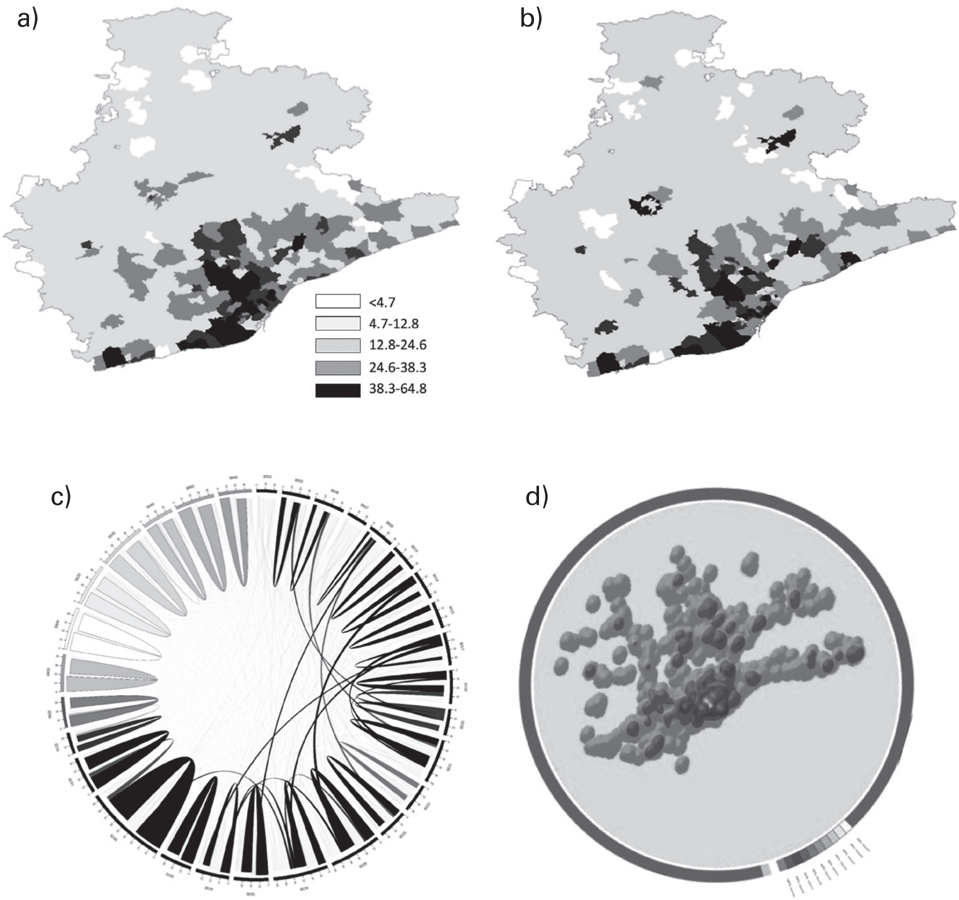

Figure 5.6
Flows of money in the city of Barcelona: (a) expenditure at the household location; (b) revenues at the retail location; (c) money flows abstracted between the household locations; (d) geolocated tweets across the whole city.

posting information about themselves and others as well as their opinions on various web sites accessible by others. The various information that is posted may not exclusively be about social issues, and an increasing proportion of social media is no longer focused on social interactions but involves information relevant to the operation of business, government, health, and education. For example, in the United Kingdom, hospital and medical appointments are invariably communicated by the National Health Service to patients through text messaging. A small proportion of this information is geotagged by users who open their devices to geo-positioning satellites (a function which invariably needs to be activated by the user), generating data that can be visualized in space and time. In terms of generic information that can cover any subject area (within certain ethical limits), only a very small proportion is geotagged, often less than 5 percent, and usually less than 2 percent. This media is generated in real time and usually absorbed by those who are exposed to it, also in real time, as in short text messages associated with services such as Twitter.

This data is also highly biased in that the young tend to be much heavier users, with many of the portals that have been set up to log and record such interactions heavily dominated by younger rather than older adults.[30] This is partly related to the ownership and use of smart devices, laptops, and tablet computers. But unlike traditional data recorded from periodic censuses—administered so that the maximum of any sample population or total population itself is captured—it is very hard to generalize new media data with respect to its meaning for entire populations. Because meaning has to be extracted from it, often using various forms of data mining, the use of much of this data is controversial, unclear, and uncertain. These are limits to the online world that currently make it highly problematic to use such data for understanding social and economic interactions that characterize the city and display its functioning.

Where such social media data has been visualized in spatial terms, it often tends to show what we know already—a repeated criticism, of course, of all intuitively acceptable results that are explicable in hindsight. Quantity of tweets, for example, tends to co-vary with the density of different populations, although it is clear that such data also co-varies with the presence of younger people in a city, who tend to be more active during the evening than the early morning. Figure 5.7 shows the distribution of tweets over a typical week for four world cities. These "new city landscapes," as Neuhaus[31] calls them, are based on extracting tweets from one of the Twitter APIs. All the geotagged tweets for the week in question are plotted within a radius of 30 km (nearly 19 mi) of a central location in each CBD. These show an obvious correlation with population density, although we do not show these densities: we simply assume the reader has some sense that these cities have typical population densities, which decline from the core to the periphery. This is central to the standard model and is consistent with all that we will say about the spread of cities in the next chapter. However, in each case,

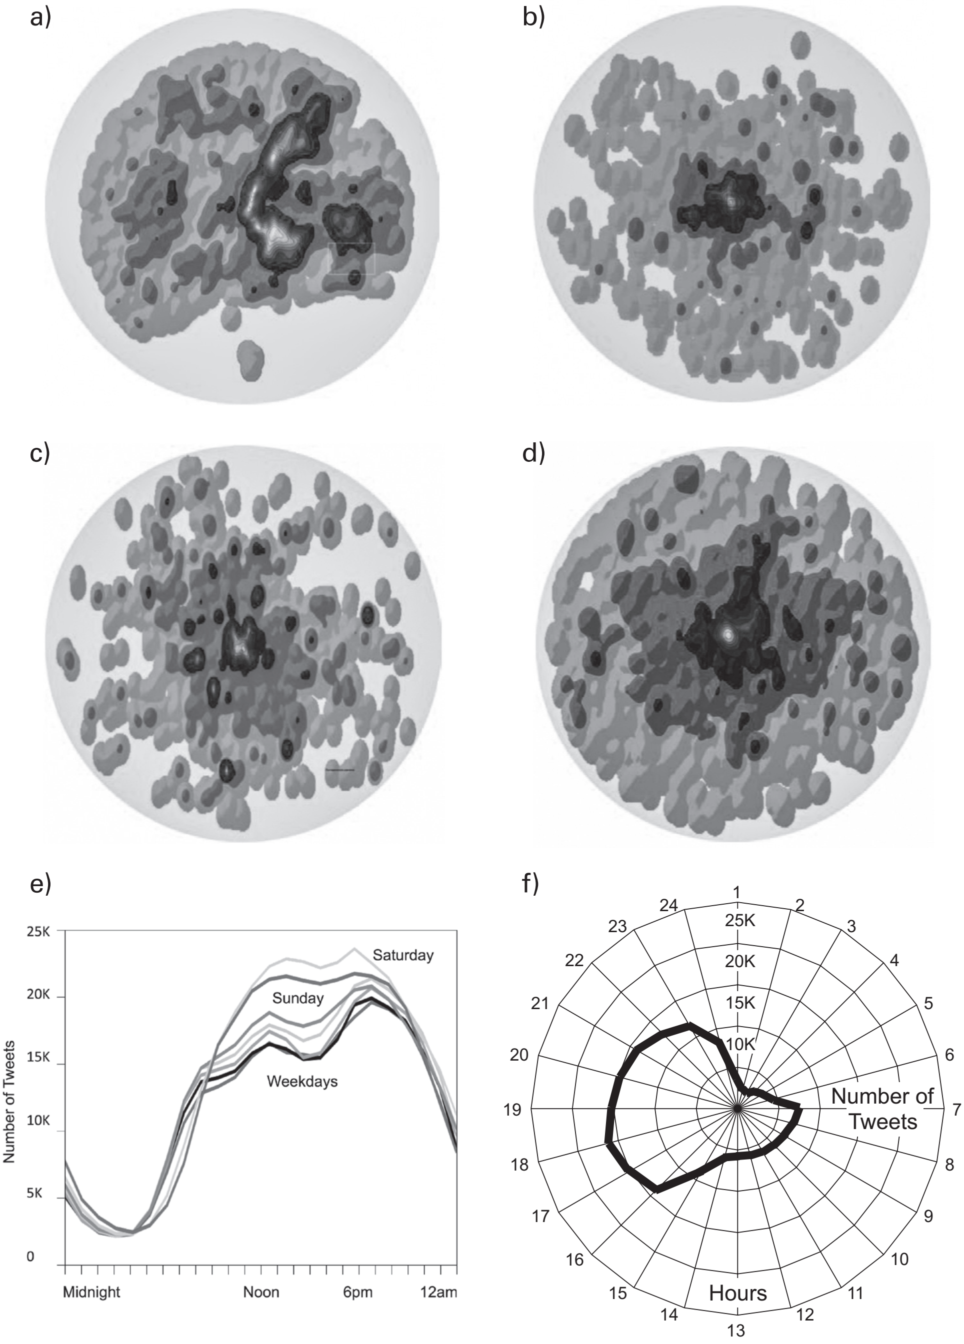

Figure 5.7
Volume of tweets in world cities: (a) New York; (b) Paris; (c) Moscow; (d) London; (e) London: daily time series; (f) London: average daily volume.

we can see differences, and we will concentrate on London in figures 5.7(d) to (f). In 5.7(d), there is much higher density to the north of the CBD in Camden, which is a retail and entertainment center. The daily profile we show in figure 5.7(e) in traditional form, and then in 5.7(f) in circle-like form. This illustrates that there is a small peak in tweets when people wake in the early morning, and a slight decline in volume as the working morning progresses, but then a rapid rise after midday that continues all the way through the afternoon, rising to a peak at about 7pm. A plateau is reached by mid-evening, which then begins a precipitous decline as people leave entertainment and other venues, preparing for sleep in the late evening.

Of course, we are guessing about what this profile means, illustrating the basic problem with all online data being routinely captured at fine temporal and spatial scales: it is impossible to tie it to independent, observable processes whose functions generate that data. We simply do not know what people are doing when they tweet, and we barely know where most of them are located. Unless there are some very dramatic changes in how we engage with media and information that captures our private behaviors, it is most unlikely we will ever be able to get much further than the illustrations we are now able to extract shown in this chapter. But inventing the future city may be as much a matter of adopting different behaviors as building new kinds of physical form that relate to the changing functions brought upon us by new technologies and changed lifestyles. In this sense, social media data such as that from Twitter potentially contains much of interest pertaining to who is contacting who about what. In short, there are a multitude of network connections embedded in such data—but the difficulty of extracting it is on a par with the difficulty of interpreting it. The best that we seem to be able to do at present is impute in some rather vague way social links based on people who retweet information targeted in some way at themselves. If a person tweeting retweets something and addresses the tweeter who wrote the first post, then this might be regarded as some sort of social link. If a tweet is sent that includes reference to somebody with a Twitter account, then this might also be regarded as a link. There may even be the possibility of extracting peoples' identities from the message itself, though this must be approached with caution.

There is enormous ambiguity in inferring such networks from data of this kind. It is not at all clear what a retweet network really means, other than the fact that some have decided the information is relevant to their contacts. The link is not necessarily reciprocal, and it probably requires an analysis of the demographics of their contacts to add any further meaning to this. This is not only relevant to Twitter, but to all social networks that involve this kind of real-time activation. If this kind of analysis is extended to sites such as Facebook, there is more content to be extracted and analyzed. It is possible social networks are then more like traditional networks, where actual contacts are measured directly. We could spend a lot longer talking about these questions, but we have to leave them unanswered. They pertain to a future that we are in the process of inventing and reinventing, and there is likely to be much more volatility in these kinds of media than in traditional physical functions, which will continue to dominate our cities in the foreseeable future.

Our last foray into this kind of data relates to data that individuals post to remote sites that record relatively unambiguous content with accurate space-time coordinates. There are now literally thousands of explorations of social media data from Twitter, Instagram, Flickr, Facebook, and so on, which have data that is encoded and temporally tagged, but little of this has been integrated in a comprehensive way. In figure 5.8(a), we show a typical map from Eric Fischer that is a visualization of tweets in London over a long period of time. At this metropolitan scale, the location of these data marks out street and railway networks. If we scale this up to the European level in figure 5.8(b), we reveal cities and urban clusters.[32] These are strongly similar to the flow maps explored in the previous chapter. In figure 5.8(c), we show maps that are extracted from the location of date- and place-stamped photographs posted to Flickr for London. Fischer has visualized both local and tourist posts from this data, to an extent showing how different kinds of posts are integrated and comparable spatially and temporally. Last but not least, we show pairs of friends globally from Facebook data visualized by Paul Butler.[33] This world map and the European map of tweets in figure 5.8(c) show the kind of global geography—or "connectography," as articulated by Parag Khanna[34]— most cities are rapidly joining. To conclude our brief foray into this digital world and the way it is changing cities and introducing new dimensions into our understanding of cities and city planning, we will conclude with a little speculation on what all this means for flux and flow in the future city.

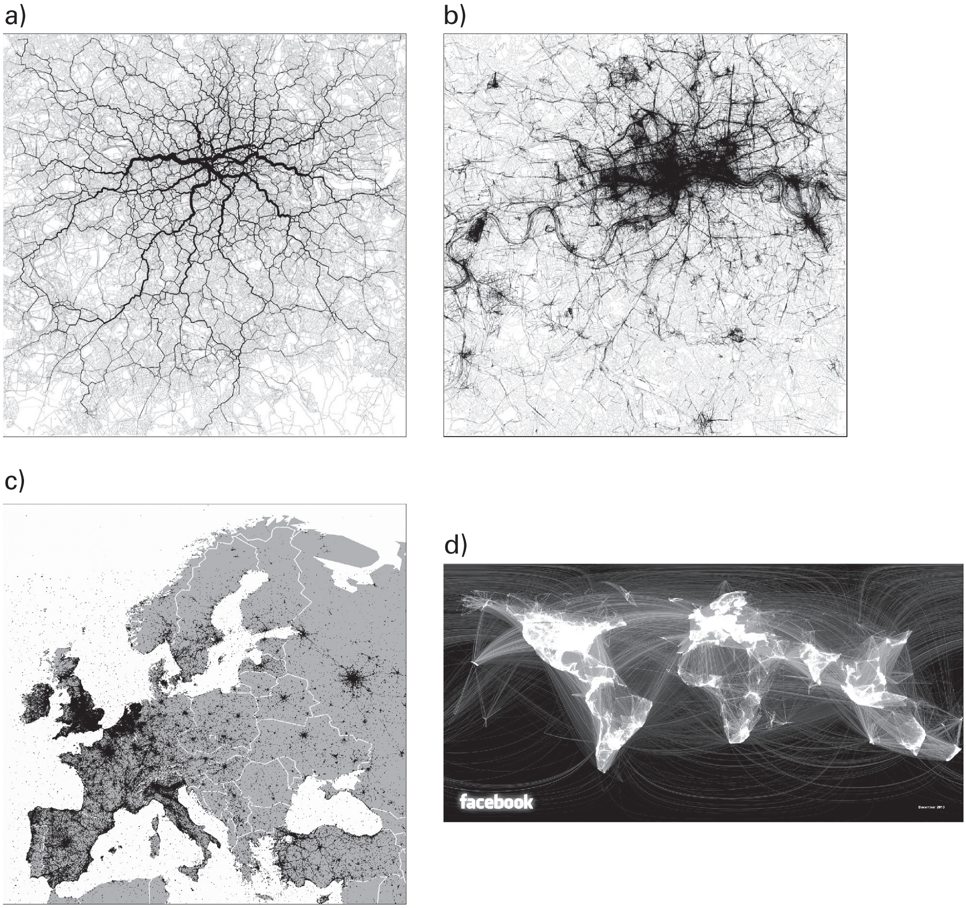

Figure 5.8
Geolocated tweets, Flickr photos, and Facebook links: (a) Fischer's London tweets; (b) Fischer's London Flickr photos; (c) Europe tweets; (d) Butler's global pairs of Facebook friends.

Future Cities, Liquid Cities

Cities connect people, but, as pointed out several times already, they do not function as machines but as organisms. People are continually connecting and disconnecting as the city grows and changes. This is very evident in all we have shown in this chapter. We have taken a tiny sample of these new kinds of digital connection that are now possible and illustrated how these vary over the day to the year. Imagine a multitude of these streams forming an enormous connected skin across the physical fabric of the city, a skin composed of many layers, continually activating and deactivating as new nodes are added to the networks, old nodes disappear, and existing nodes transform their roles. Imagine many different kinds of physical and digital infrastructure supporting these networks and changing at different rates with respect to the data pushed through their structures. The rational for such exchange is both social and economic, with different networks being used for one or the other or any combination of functions associated with how the city works as a social and economic organism. The challenge is not simply to represent such functions and channels, but to integrate them using some model of how the whole system works. Both the operation of the system, its maintenance and its quest toward optimality for its users, whoever and whatever they represent, requires a powerful theoretical framework for exploration and conditional prediction. So far, a century of effort has only produced the simplest of toy models of how urban economies and ecologies work. These are still far from the mark. The ever-growing global systems of cities is becoming more complex daily, and our existing theories, introduced albeit very crudely here, are woefully inadequate. In fact, we are running to stand still. The challenge is to make sense of all this, and quickly.

Should we be surprised by this dilemma? Probably not, because urban evolution and the cities that define it are continually changing their form and function anyway. Invention is what this is all about. Cities have been getting more complex exponentially since the dawn of history, as Jane Jacobs so persuasively argued, since tiny pockets of urbanization came to be established as far back as the Stone Age.[35] Add to this changes in technology, which in turn are part and parcel of the demographic explosion, and it is clear that we have never been able to predict what comes next. In that sense, all that actually takes place is invented by ourselves. In short, we cannot predict our inventions, and thus it is hard to know how the new digital structure of cities will play out in the next 100 years and beyond. We are all likely to be connected in a giant urban cluster by the end of this century, which will be the physical manifestation of an intricate global nervous system created by the world of information technologies we have attempted to picture here. When information is instantly available, from anywhere and anybody, at any time, things in cities begin to look different from their historic past. Cities are becoming considerably faster in their responses to new information, to innovation, to physical change. Populations with more information are able to make decisions ever more coherently and at faster and faster rates. In this sense, the city manifests a new liquidity of action, a confluence of light and speed, and ultimately becomes a "liquid city": a place where physical desires, face-to-face contacts, and digital deliberations provide a new nexus of innovation. Flows, networks, and connections, rather than inert structures, dominate this physicality as infrastructure comes to represent this new liquidity built on layer upon layer of flux and flow.

As we have shown in this and the last chapter, the exemplar for understanding and representing the city is the network. The various examples shown here illustrate network structure at many levels, where hierarchy is implicit. This kind of structure reflects growth around a series of market centers and other hubs or cores in the urban landscape, but these hubs are changing as more and more human activities come to occupy the ethereal rather than just the physical space. Networks spread out to capture consumers and producers, who come together to buy and sell in markets, while the traditional function of transportation in cities is to connect people in economic exchange. Flows of people to work and to markets to produce or consume through trade and exchange are the outward manifestations of the most obvious flows of energy forming the cement that keeps the components of the city together. Networks to channel these flows cannot develop everywhere, so they spread out, tree-like. as if reaching out into the air or the soil, the way a tree grows out in search of energy to sustain it in the most economical way. To an extent, this is what Alexander argued many years ago[36] and what the idea of the city as a fractal geometry is all about.[37] The capacity of these networks increases according to the population that can be sustained by each of its nodes, and when a node reaches capacity, new ones appear, as in edge cities that are rapidly changing the dominant form of the contemporary city from monocentricity to polycentricity. The way cities fill their space is intimately related to the way we try to use space efficiently, in two and three dimensions, and the form that results is the product of many decisions that grow the urban fabric from the bottom up. We will explore this kind of space-filling in the next chapter and speculate on how this physicality is extending into the virtual world, which increasingly influences the way we behave in the physical city.

Cities are becoming their own sensors at the most elemental level, as their physical fabric is being automated in ways that enable us to monitor their performance and use. But these combined material and digital forms are also being overlain with digital skins that seek to enable populations to use the city in countless different ways, such as figuring out in real time what services are located in different places, where friends and acquaintances reside, and what physical means of transportation there are to move to distant locations. When we combine this with more basic sensors that reflect the way buildings are working, we are augmenting our reality in ways that have not hitherto been possible. At higher spatial and temporal scales, our ability to sense the city remotely offers new insights into how cities are growing or declining, how they are being used over longer periods than real time, and how we might identify problems that emerge at a global scale more associated with individuals operating in real time. All of this offers enormous opportunities for a new age of urban design that takes account of the city as a self-organizing system blending the physical with the digital and offering new opportunities for social and economic interaction. Much of what will occur in this century will disrupt the traditional city in ways we have not yet invented. In the rest of this book, our preoccupation will be on the charting the pathways along which this might happen.