

# From Plan-Demanded Data to Plan-Ready Information: A Rationale for Comprehensive Urban Knowledge Infrastructures

Fabio Carrera and Lorlene Hoyt

**B**ASED on our professional experience in urban planning in the United States, Italy, the United Kingdom, and India, we recognize that cities often lack a comprehensive “knowledge infrastructure” on which to base planning decisions. Instead, what we discovered to be the most prevalent tendency for the various departments of a modern city is a form of “ad hoc-ism” whereby data are collected for specific purposes and then quickly forgotten or stored in inaccessible places. Although some systematic data collection takes place—mostly for regulatory or revenue-generating purposes such as for permits, licenses, and property assessments—even these data are generally hard to obtain or utilize, quite often due to real or imaginary privacy concerns. Frequently, access to important information is made possible only through personal connections and by means of under-the-counter transfers which bypass official channels. Moreover, we have also witnessed instances where a single department hired private-sector consultants to collect the same data multiple times, and situations whereby the same data were collected by different departments simultaneously. In short, redundancy and waste seem to be endemic when it comes to municipal data collection.

Evans and Ferreira, Jr.  
Ferreira, Jr.  
Hoyt et al.

---

*Journal of Urban Technology*, Volume 13, Number 2, pages 3–23.

Copyright © 2006 by The Society of Urban Technology.

All rights of reproduction in any form reserved.

ISSN: 1063-0732 paper/ISSN: 1466-1853 online

DOI: 10.1080/10630730600871999

Urban professionals who are engaged in maintenance, management, or planning activities use information daily and have been receptive to the adoption of computers to organize municipal information. Since the first commercial Geographic Information System (GIS) appeared in the mid-1980s, there has been a steady increase in the use of GIS at the municipal level of government. The diminishing cost of computer hardware and software has led to a proliferation of homegrown GIS initiatives to address the specific needs of municipalities. Increasingly, there have been attempts to harness the richness and diversity of such independent activities to reduce wasteful redundancy and to maximize the synergistic potential of a coordinated approach to geospatial information management.

Top-down initiatives emanating from the national level have led to the creation of Spatial Data Infrastructures (SDIs) which, in turn, frequently include provisions for a core set of so-called "framework data." Simultaneously, bottom-up efforts focus on the role of neighborhoods in the development of a fine-grained spatial data infrastructure through Community Statistical Systems (CSS) and the like. Citizen groups, often by means of university-community partnerships, are producing the neighborhood-level equivalent of national framework data, through such efforts as the National Neighborhood Indicator Partnership (NNIP). Such efforts consolidate indicators of urban well-being using public administrative data sources. Meanwhile, technical standards that allow the exchange of spatial data are also being developed primarily by the Federal Geographic Data Committee (FGDC).

Despite this positive ferment, planning professionals approach the development of Planning Support Systems (PSS) in a laissez-faire manner. The explanations for this lackadaisical attitude are many and include: their inability to take full advantage of the technology; organizational, institutional, and sociocultural issues; and their preoccupation with gathering useful data for the plan at hand which at best results in the mere computerization of manual tasks. In fact, with the exception of a small number of well-funded and established comprehensive top-down efforts, progress has been slow and "the application of PSS is currently still in its infancy." Another major stumbling block is the quality and availability of data, where issues of privacy and intellectual property further hinder the development of GIS enterprise systems. Yet, our experience is that data protection obstacles can be overcome, especially when it comes to data that are to be used internally only, within a single city department.

Nedović-Budić and Pinto 2000

Masser and Wilson  
Budić

Klosterman

National Research Council

Nedović-Budić et al.  
Masser and Wilson

Tulloch and Fuld

Talen

Sawicki and Flynn

Coulton et al.

Budić

Innes and Simpson

Nedović-Budić

Geertman and Stillwell 2004

Innes and Simpson

Campagna and Deplano

Nedović-Budić

Masser and Wilson

Zuboff

Arun and Yap

Wilmersdorf

Geertman and Stillwell 2004

Yeh

Carrera and Hewitt

Planners are voracious consumers of information, but they rarely produce new information themselves. Comprehensive systems to bring together multi-purpose systems for second-order spatial analyses are needed but not available. Current trends indicate a move toward the development of local geographic information strategies to capture the fine-grained urban data that community statistical systems require. A discussion about the importance of supporting the development and maintenance of local databases is also emerging, and the task of constructing distributed municipal information systems from a series of networked systems that are connected via the World Wide Web and developed in a coordinated manner is a potential reality. According to leading experts, a first challenge “lies in striking a balance in the degree of centralization of data storage, administration, and procedural control while serving the needs of the community [. . .].” Additionally, the University Consortium for Geographic Information Science (UCGIS) declares,

Innes and Simpson  
Brail and Klosterman  
Geertman and Stillwell 2003  
Geertman and Stillwell 2004  
Tulloch and Fuld

Nedović-Budić

Kelly and Tuxen  
Keating et al.  
Nedović-Budić and Pinto 1999  
Nedović-Budić and Pinto 2000

Keating et al.

As the variety of geospatial information and data resources increases each year, the demand for understanding and building sustainable information and knowledge structures remains a critical research challenge for the geo-spatial information community.

Shuler

So the problem today is not the availability or capability of technology for planning, but rather the availability of “good” fine-grained, up-to-date data. The other missing piece is the creation of systematic storehouses for urban knowledge. One way forward is to embed—in the planning community—an appreciation for the value and importance of information at the local level with regard to urban maintenance, management, and planning. This philosophical shift would enable a sea-change to take place in how cities collect and organize information. It seems possible to envision being able to gather, organize, maintain, use, and re-use the datasets that would feed a comprehensive urban knowledge infrastructure. It remains a tall task, but it is no longer an insurmountable one. The transaction costs and complexities associated with geospatial data collection continue to decline and the impacts of technological and organizational change have been understood and can be factored into any economic calculation of benefits versus costs. In short, we know how we could assemble the system and we may be able to afford it.

Budić

Innes

Evans and Ferreira, Jr.  
Nedović-Budić and Pinto 2000

What remains elusive is how to put together and finance a well-oiled machinery that will keep all of the datasets organized

and up-to-date so that ever-improving applications can run each aspect of a municipal operation in an efficient and cost-effective manner. Fortunately, a majority of the characteristics that make up the physical city change very slowly, if at all, and are thus amenable to a gradual and systematic collection effort, the bulk of which would only have to be conducted once. Until now, the apparent complexity of collecting and organizing such a multi-dimensional body of information has discouraged a wholesale approach to the accumulation of municipal information.

Carrera 2004

This paper argues that municipal planners should spearhead a move toward this end. The marginal returns one can obtain by systematically collecting and archiving fine-grained urban data are beginning to outweigh the transaction costs that such collection efforts would entail. The case studies we present herein showcase the first- and second-order returns that we were able to extract from our progressively accumulated knowledge base when we used our data for a specific, immediate purpose and then later reused these same data for new research on a different topic. These cases demonstrate the obvious gains attainable by automation as well as the unforeseeable advantages one achieves through the development of a comprehensive urban knowledge infrastructure.

Central to our argument is the notion that urban information is a resource that should be maintained like any other city-owned system. Investments in municipal information systems should, therefore, be viewed as capital outlays and steps should be taken to ensure that taxpayer money is well spent. In the same way that a municipality would not consider rebuilding its sewer system over and over again, we argue that information about important urban elements should not be lost, underutilized, or repeatedly collected. Drawing on our collective professional experience, we conclude that planners are perfectly positioned to play a role in transforming data to knowledge and should serve as catalysts for a long-overdue transition from hunters-and-gatherers of “plan-demanded” data to farmers of “plan-ready” urban information, thus promoting the institutionalization of a comprehensive urban knowledge infrastructure.

### **From Data to Action**

City data, for the purposes of this paper, include all data that are relevant to municipal maintenance, management, or planning. Indeed, the planning process is predicated on the availability

of data, but information is rarely available. Urban planning today is largely based on ad-hoc data collection, meaning it is gathered from a range of agencies that are administratively isolated and concerned with different issues and integrated as needed. The term we use to describe this mode of operation is “plan-demanded.”

Ferreira, Jr.

Automation plays a certain role in this process, in that some planning data are collected fairly rigorously by some government agencies, but the tendency toward automation in this field has been limited, for the most part, to areas that are under strict regulatory control (e.g., land use) or that generate municipal revenue (e.g., parcel ownership). Recordkeeping in such instances has always been necessary to the proper functioning of civil society, so the introduction of information technologies has been merely a convenient way to make the process faster and smoother. Generally speaking, the representation of space in many municipal computerization efforts has been shortchanged. At best, locations are represented by address, with all of the standardization and referencing problems that such an approach entails. A systematic approach to the acquisition of fine-grained city knowledge is still considered too cumbersome, even after the introduction of the first geographic information systems in the 1980s.

Knowledge is supported by—though not exclusively composed of—“hard facts.” The “hard facts” are, unfortunately, not as available as one would imagine them to be. In fact, many distinguished planners of the past as well as many contemporary observers of urban affairs clearly point out that we are not doing a really good job of knowing our cities. Despite the relative permanence and immutability of the physical elements composing our urban realms, knowledge of our cities is not as developed as we would like it to be. Although data are gathered daily for a variety of reasons, information is not necessarily obtained as a consequence, and knowledge is therefore hardly augmented in the process. The seemingly subtle differences between data, information, and knowledge are quite apparent in the fields of urban maintenance, management, and planning. Data are all too frequently collected to satisfy very specific needs; they are mostly treated as mere documentation and, therefore, are rarely organized into information that can be used for other purposes. Thus, seldom do they contribute to the creation of knowledge on which decisions and actions can be fruitfully based. Planners, by the nature of their trade have to learn how to collect, use, and share knowledge using a variety of rational or formal rhetorical forms. Our collective

Olmsted, Jr.

Geertman and Stillwell 2004

Mandelbaum et al.

praxis, straddling the worlds of professional planning and academia, leads us to believe that a gradual accrual of fine-grained municipal information is feasible with today's technology and could be realized, one department at a time, through the advocacy of city and town planners in their roles as primary beneficiaries of "plan-ready" information. The sections that follow illustrate the benefits of transforming "plan-demanded" data to "plan-ready" information using some examples from the direct experience of the study of the inner canals of Venice, Italy by one of the authors.

Carrera 1996

### **Plan-Demanded Data Collection**

Examples of "plan-demanded" data gathering abound. In fact, most data-gathering outside of regulatory or revenue-generating operations usually fits into this category. The authors have participated in a range of "plan-demanded" data collection campaigns in several national contexts. For instance, Carrera has supervised several teams of students from the Worcester Polytechnic Institute (WPI) in the collection of data throughout the City of Venice. These data are currently used by Insula S.P.A., a public-private entity in charge of the maintenance and restoration of the Venetian waterways. The data collected included: measurements of the physical dimensions of the canals, including the water depth and sediment levels at the bottom; a catalog of all sewer outlets and wall damage along canal banks; measures of the water currents in the canals; counts of the boat traffic in the canals and quantification of the wakes produced by passing motorboats; an inventory of all bridges spanning the canals and an assessment of their state of repair; a census of all boat docks and their condition; and a series of studies to quantify the amount of cargo delivered to each island in the city. These data were later used by the sponsoring agencies to carry out specific actions related to urban maintenance, management, and planning that required immediate attention. Below, we present some specific examples of "plan-demanded" data gathering.

The first example of "plan-demanded" data collection concerns the measurement of the depths of Venetian canals. Such measurements were originally intended to provide a useful quantification of the amount of sediment deposited on the bottom of canals that would allow an estimation of the dredging efforts required to re-establish adequate drafts that would permit emergency boats to travel the canals unimpeded. Our bathymetric dataset included 7,768 individual measurements along 850 section

lines on 130 canal segments. Approximately 55 soundings were performed in each segment, at one-meter intervals across each of the section lines. The average of all the depth measurements in one segment was used as the overall average depth of the entire segment. (See Figure 1.)

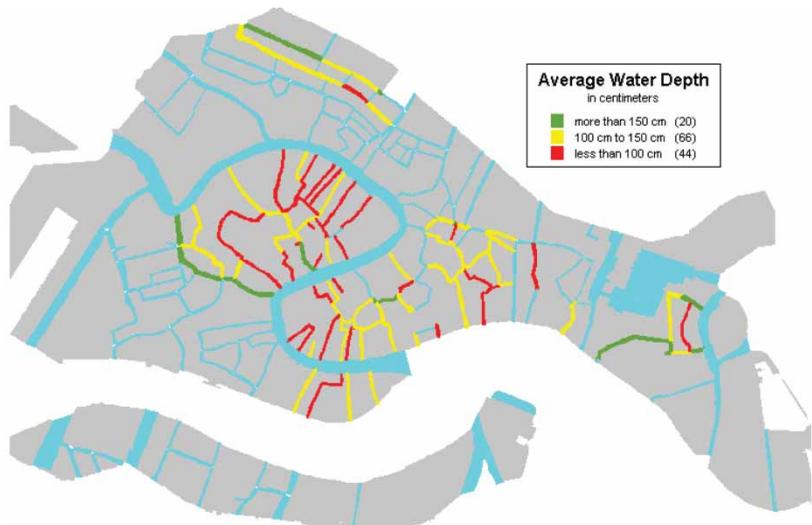
In order to make the average depth as realistic as possible, the average was computed omitting measurements along the canal walls on both sides of each section. Since we were the first to systematically determine these dimensions, we were also the first to produce reliable estimates of the volumes of sediment that needed to be removed from each canal segment.

Determining the consequences of sediment buildup on the navigability of each canal segment was also a challenging task. Average depths are only marginally useful in this context. Depending on the draft of one's boat and the tide conditions, one might opt to avoid canals with low average depths, but there would be no guarantee of being able to navigate the segments with medium average depths or, paradoxically, even the ones with the highest average depths. One shallow point in an otherwise deeper channel would be sufficient to completely impede passage, even though that canal segment may be very deep on average as a whole. Alternatively,

---

**FIGURE 1**  
**Average Water Depth in the Venetian Canals**

---



---

Source: Carrera 2004

---

one may be tempted to identify the single shallowest data point in a canal segment as the navigational bottleneck where boats may run aground, using the 7,768 individual depth soundings at our disposal. However, this approach would also be fallacious, since the absolute shallowest points are almost always along one of the canal walls. Knowing that such points are the shallowest would be useless, since no boat could possibly travel along the canal walls. In fact, a typical boater would attempt to tread a route through the deepest part of the canal, so that, at any moment, the boat would be passing over the deepest point of a particular cross-section, staying well clear of the shallowest. Therefore, what would be truly useful is to know where these deepest points are across any section and then identify the shallowest one of such deep points.

So, we did precisely that using a two-step manipulation of the fundamental soundings dataset. First of all, we identified the deepest points of each section, and then we selected the shallowest of these to truly determine the navigability bottlenecks. In our research, we defined what we call a “navigability axis” as the line connecting the deepest points of each of the bathymetric sections. Thus, even though we could not vouch for the depth of any water between sections, we could determine the point along this navigation axis where the water was shallowest, hence providing a more useful indicator of navigability to local boaters. This “shallowest of the deepest” concept is one way to think about the distinction between data and information.

This example reveals the usefulness of collecting and retaining the fine-grained bathymetric soundings instead of “throwing them away” once the average depth has been calculated. More importantly, this case demonstrates how the same fine-grained “plan-demanded” data regarding the depth of Venetian canals were not only of immediate use for the canal-dredging plans being developed by *Insula* (first-order), but were also central to the success of more sophisticated analyses (second-order) such as the impact of sediment accumulation on canal navigability.

Our second example of “plan-demanded” data collection pertains to the bridges of Venice that we cataloged to assist *Insula* with its maintenance activities. In the 1990s, we began to inventory the 472 bridges that span the canals of Venice. Our inventory included information about the surface area and material of the bridge; the thickness of the structure (for determining the space available to run pipes and cables through the bridge to join two islands); and the height of the span (to determine the clearance available to boats during high tides). (See Figure 2.)

Despite the fact that bridges had been there for centuries, no one had ever attempted a systematic inventory of the physical characteristics of all bridges in the city. We also surveyed the conditions of a variety of elements of each bridge to determine the overall state of conservation, which had a direct impact on the scheduling of maintenance interventions. Once again, this “plan-demanded” inventory produced instantly useful data on both the permanent features as well as the dynamic attributes of the bridges.

The third and final example of “plan-demanded” data collection in Venice involves boat traffic in the canals. These data were needed by the Public Services Department, the agency that oversees traffic management and regulations. Row boats and sail boats have traveled the canals of Venice since the fall of the Roman Empire, but motorboats appeared in the middle of the twentieth century. Motorboats have since become one of the major problems the city is facing today. Boat traffic in Venice not only generates congestion, but it is a more insidious phenomenon than its land cousin. Unlike automobiles, motor boats can physically destroy the foundations of buildings along their paths, due to the water turbulence and wakes that they create. (See Figure 3.)

**FIGURE 2**  
**The Information System for Bridge Maintenance**

**Informazioni Generali**

Ultimo Restauro  
 Data Ultimo Restauro  
 Anni trascorsi  
 Rilevamento Disegni

**Informazioni Generali**

Nome Ponte **Ponte Carlo Goldoni**  
 Codice Ponte GOLDO  
 Anno di Costruzione  
 Rio Attraversato Rio dei Fusari

**Informazioni Architettoniche**

N. Campate 1  
 N. Rampe 2  
 N. Gradini Nord 14  
 N. Gradini Sud 14  
 Sommità  
 Larghezza 4.64 m  
 Sporgenza Nord 0.4 m  
 Sporgenza Sud 0.7 m  
 Luce Effettiva 12.2 m

**Decorazioni Artistiche**

Decorazioni  
 Condizioni Decorazioni Buone

**Spallette**

Altezza Media 96 cm  
 Spessori  
 Spessore 63 | 53 | 43 | 53 | 60 cm  
 Spessore Minimo 49 cm

**Materiali e Condizioni**

Materiali	Condizioni
Pavimentazione Gradini Trachite	Buone
Pavimentazione Sommità Trachite	Buone
Alzate Pietra d'Italia	Medie
Spallette Metallo	Buone

Condizioni GENERALI **Buone**

Source: Carrera 2004

---

**FIGURE 3**  
**Traffic Congestion in Venice**

---



*Source: Carrera 2004*

---

Despite the massive exodus of Venetians away from their city of birth, boat traffic continues to increase in order to cater to the needs of a growing tourist industry. According to our own calculations, traffic has almost doubled in the last 25 years, while the total population has declined by 50 percent in the same period. People who live along the canals are watching their dwellings crumble and many property owners have begun a public protest against the daily assault on their homes. This problem is particularly intense along the primary arteries. Despite the apparent intricacy of the web of canals, the entire water network can be schematically simplified to just a few primary routes where most of the boat traffic concentrates; there are 367 segments and 182 inner canals in the city, yet there are only a handful of main thoroughfares. The relative levels of traffic in each of these thoroughfares were not systematically quantified until we began to record traffic flows. Between 1992 and 1994, we recorded almost 60,000 transits in almost 400 hours of counting as part of an effort to which 21 WPI students dedicated a total of about 15,000 thesis-hours. Our traffic counts were of immediate use to the City of Venice, and the methodology that we developed for conducting, archiving, and analyzing the counts has been officially adopted by the city in all subsequent traffic campaigns in Venice. The traffic data, as well as the aforementioned bathymetries and bridge datasets, have become integral parts of the growing body of “plan-ready” city knowledge that we were able to re-utilize in other contexts such as the one we present in the next section.

Fiorin and Carrera

Carrera 1996

## Plan-Ready Urban Information

While involved in the aforementioned “plan-demanded” projects, we immediately realized that such efforts would be much more effective if they not only contributed to the pressing needs of the agencies that commissioned the studies, but also contributed to the long-term creation of a knowledge infrastructure that could be re-used in other contexts or for other purposes. While Insula could use the canal data for its immediate needs, permanent and immutable features such as the canal lengths and widths could certainly come in handy for some other purpose at a later date. More importantly, the canal coding scheme that we developed, which assigned unique identifiers to each segment of the water network, would undoubtedly be useful for posterity. If all future data-gathering used the same spatial infrastructure scheme, it would be possible to compare and correlate datasets referring to the same canal segment at any time. Accordingly, when we collected our data on all of the various elements of the urban realm, we always did so with an eye to this fundamental infrastructure of knowledge that could be reusable by other researchers or government agencies.

To institutionalize the processes that led us to our own “plan-ready” intuitions, we are proposing that planners actively engage in promoting a space-based representation of the urban realm based on the fundamental, quasi-permanent physical elements that are already the object of regular municipal attention for maintenance or management. While this may not be a novel idea in itself, the innovation we are suggesting would lie primarily in the manner in which these data could be systematically collected, and especially updated, by capturing transaction data and even some low- to no-cost snapshots, starting from a few key areas that are especially relevant to planning. An important aspect of our approach is to focus first and foremost on the permanent features of the urban world; once recorded and organized, these data would require very little upkeep, thus eliminating any redundant effort to re-collect them in the future for different purposes.

The city departments that would benefit most from a structured approach to the representation and computerization of the urban features that are under their jurisdiction would take the lead on such an endeavor. It is at this level that the systematic approach we propose can be most effectively overlaid on ordinary municipal operations where the tradeoffs between maintenance necessities and the added requirements of the encoding of city knowledge are most advantageous. In short, this approach

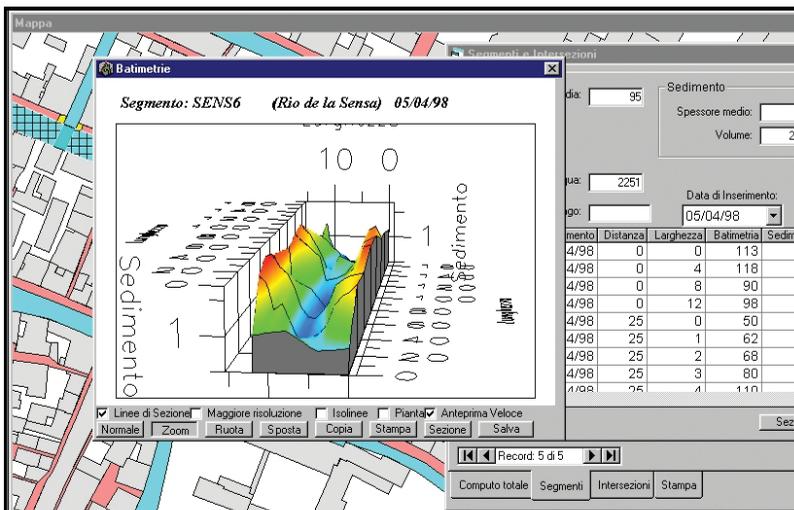
promises to produce “plan-ready” information and urban planners would greatly benefit from its existence.

In fact, each of the “plan-demanded” data collection cases covered earlier yielded re-usable “plan-ready” information that was encapsulated in a custom-designed, multi-media information system designed to support municipal maintenance, management, and planning efforts. For example, to support dredging operations, we co-designed a stand-alone application that allowed the visualization of our bathymetric data and produced three-dimensional interpolated visualizations of sediment accumulation. (See Figure 4.)

To facilitate the maintenance of bridges, we produced an application that integrated the data collected about each bridge, as well as photographs of ramps and arches. A separate screen displayed the state of conservation for different elements of the bridge and calculated an overall condition score that prioritized restoration interventions. Finally, we designed an interface to display information regarding temporary canal closures and other traffic management decisions.

While the systems described above were primarily geared at making data available in a compact and complete manner, the

**FIGURE 4**  
**The Information System for Managing Canal Maintenance**



Source: Carrera 2004

information regarding bathymetries, bridges, and traffic can be used to produce more sophisticated, second-order applications for the purpose of conducting advanced analyses, simulating future trends, or identifying causal relationships. For example, Carrera contributed to the development of a traffic model for Venice that allows city officials to forecast the effects of canal closures and to simulate the consequences of regulatory changes such as the institution of new one-way canals.

The next example clearly demonstrates how “plan-ready” information can free planners from the tedium of redundant and inefficient “plan-demanded” data collection and allow them to concentrate on more beneficial high-order tasks. “Plan-ready” datasets cannot only assist planners and municipal staffers with day-to-day municipal maintenance and management activities, but they also provide the requisite information to support more advanced analyses and projects that are aimed at improving the quality of life for a broad constituency.

Caporale et al.

In the late 1990s, Carrera was asked to develop an ambulance dispatching system for the Venice General Hospital. Ambulances (which are boats in Venice) need to contend with typical transport-related obstacles such as traffic congestion, but also have to deal with the vagaries of tidal fluctuations which may make some routes impassable, either because of high tides—which make some bridges too low to pass under—or because of low tides—which make it impossible to navigate where the sediment build up has made the canals too shallow. The goal was to develop an application that would be capable of determining the shortest route that an ambulance boat would need to take to reach an accident location in Venice. To tackle this problem, we used a software package called TransCAD which figures out the shortest route (by time or by distance) and is capable of taking delays into account. Thanks to our extensive knowledge of the canal system, we were able to reuse our information about the depth of canals, the height of bridges, and the level of traffic congestion to insert appropriate delays along the routes. Unlike their land counterparts, boat ambulance drivers face a double bind: they need to first select the best water route to get close to the accident and then they need to dock the boat and navigate the best course to the emergency location over land. (See Figures 5 and 6.)

Implementing a full-fledged ambulance dispatching system in Venice would therefore require the following information:

---

**FIGURE 5**  
**The Best Water Route for an Ambulance Boat**

---



*Source: Carrera 2004*

---

1. The exact address locations—to allow the dispatcher to know the destination of the emergency boat
2. A complete network graph for the waterways—to calculate the best water courses to the destination
3. The shallowest bottleneck in each canal segment—to determine whether emergency boats may run aground at low tides
4. The clearance of all bridges—to assess whether ambulances would hit a bridge with high tides
5. The level of traffic in each segment—to estimate the congestion an ambulance may encounter on each canal on different days, at the specific time of an emergency
6. The usability of all boat docks—to decide where to tie the boat depending on the current tide levels
7. A complete network graph of the pedestrian streets—to determine the best land paths to the destination once the ambulance crew disembarks from the boat.

---

**FIGURE 6**  
**The Best Land Route from a Dock to an Accident Location**

---



*Source: Carrera 2004*

---

Typically, each of these information requirements would have entailed a massive “plan-demanded” data collection campaign and the absence of any one of them would have hindered the success of the ambulance-dispatching project. Our team, however, was able to complete the task in just seven weeks without having to collect any data, thanks to the existing storehouse of city knowledge that we had already accumulated with previous projects. Additionally, the City of Venice had already mapped individual addresses in Venice by attaching a small door-centerline to each street doorway and Carrera had previously supervised an inventory of the dock locations. As described in the earlier cases, we had captured all the canal segment centerlines, pre-determined the navigability axis of each segment, and inventoried the bridges. Thus, the task we faced was simply to associate appropriate delays for low tides and high tides when these interfered with navigation in shallow canals or under low bridges, respectively. Similarly, we inserted delays associated with traffic volumes on each segment, depending on the day of the week and the time of day of the emergency. Taking delays into account, our system produced the best water and land routes to the destination based on total travel time.

This project exemplifies the power of “plan-ready” information. In this case, we reused data collected from at least four previous projects, all of which were completed long before the ambulance-dispatch study was envisioned. What made data re-use possible was the fact that each dataset was connected to spatial features of the canal network through standardized reference identifiers (the canal segment codes).

### **The Role of Planning Professionals**

With the examples from Venice, Italy, we have demonstrated the immediate benefits that can be derived from transforming “plan-demanded” data into “plan-ready” information. We also believe that we have made the case that planners stand to gain a lot from such a transition. That is, they will be able to bypass time-consuming, data-gathering efforts in favor of more challenging—yet rewarding—high-order analyses. In these closing paragraphs, we discuss the role that planners can play in bringing about “plan-ready” municipal information systems that leverage current technologies by initiating a profound revolution in the way that cities treat information on a daily basis.

We propose that professional planners, who are uniquely positioned to act as catalysts in the conversion of “plan-demanded” data into “plan-ready” information become the champions of a civic cause aimed at systematically accumulating a comprehensive body of knowledge that will satisfy all of the maintenance, management, and planning needs of each city and town. To achieve this goal, we suggest that planning professionals begin by promoting the development of department-level information storehouses. To advance the development of such storehouses, planners should play an active role in demonstrating to the various branches of municipal government how a complete and up-to-date information infrastructure would bring measurable, immediate, and direct benefits to the department’s operations. Once planners have successfully lobbied for distributed departmental information repositories, it will be up to the individual departments to gradually and systematically capture all of the information that they need to support their operations. For instance, planners could demonstrate to the local public works department (PWD) that a thorough inventory of all roads would be extremely beneficial; the PWD would profit by knowing the exact dimensions of all paved and unpaved roads. With such information, road repairs and plowing routes could be planned and their respective costs estimated, thus providing first-order benefits for typical PWD operations. In the future, planners would also garner substantial gains as second-hand users of the same information by, for instance, analyzing the existing road widths and grades in order to establish reasonable guidelines for the conversion of private ways into public roads. The first-order paybacks, obtained by front-line municipal offices in charge of day-to-day maintenance and management functions, would be compounded by the second-order benefits resulting from the re-utilization of the same datasets for planning purposes or by additional synergies with other departments that could also reuse the data in the course of their operations. For example, the town’s traffic engineers could also make use of the PWD data to evaluate the transportation capacity of a roadway and thus enact appropriate traffic control measures.

To foster these interdepartmental synergies, City Lab at WPI is developing a Web-based urban information tool, called LOUIS (Local Online Urban Information System), which is an open-source, java-based GIS. LOUIS demonstrates how spatial tools can be used to index and correlate a number of datasets—even if they reside on different servers—using a common underlying geographic reference platform. In tandem, the Department of

Brookings Institution

Urban Studies and Planning at MIT is developing software tools called façades which will enable this approach to be adopted by various municipal departments.

## Conclusions

The concepts that we outline herein are not revolutionary, but they were simply not cost-effective for municipal departments until very recently. The falling cost and the increased capabilities of hardware and software have made a fine-grained, atomized approach to city knowledge affordable and feasible. Although there will be coordination and synchronization costs, they will ultimately be offset by the benefits of this approach.

In Venice, we have demonstrated that sizeable components of the urban realm can be systematically and exhaustively collected. Our inventories have shown adequate resiliency to change, as we have migrated them through several generations of software and hardware tools. They have also shown flexibility and re-usability as demonstrated by our “plan-ready” applications. We have also shown that the concepts developed in Venice can be exported to other cities, as demonstrated by the considerable success we have achieved in promoting meticulous and comprehensive inventories of the City of Cambridge’s curb-side parking regulations and parking meters, the City of Quincy’s public buildings, and the City of Boston’s parking facilities.

Cullen et al.  
Flynn et al.  
Blizard et al.  
Allard et al.

The systems that are most similar to our approach, in terms of their distributed and emergent nature, are the Digital Earth effort and the Global Earth Observation System of Systems (GEOSS). Though these are focused on Earth sciences at the planetary scale, they reflect the main tenets of city knowledge. They are similar in that they require distributed cooperation between independent agencies, but they have a very different—much larger—grain. Similarly, connected activities such as the Geospatial One-Stop and the National Map are producing appreciable results with many local initiatives being spawned monthly. Bottom-up initiatives, such as neighborhood information systems and the National Neighborhood Indicators Partnership (NNIP) also resemble our distributed approach, but focus on socioeconomic indicators rather than physical elements that municipalities maintain, manage, and improve.

Carrera 2004

Hoyt

Our approach is a hybrid that combines the emergent and gradual deployment of data infrastructures adopted by Digital

Earth and the National Spatial Data Infrastructure (NSDI) with the bottom-up nature of neighborhood information systems. Our focus is at the municipal level, and more specifically, we are interested in changes within departmental offices. We think that our approach, now that it is technically and economically feasible, promises to be more sustainable since it leverages the self-interest of front-line municipal departments. Department heads with an interest in moving toward the paradigm we propose should begin by creating a municipal spatial framework in which to plug urban data as they are collected. The thorny issues of coordination of distributed agents, synchronization, and replication that are standard fare in the field of information technologies will eventually be resolved; however, bottom-up, high-resolution data collection can be started at any time. For those cities and towns that are already collecting and mapping urban data, the change would be toward data of a finer grain, richer attribute sets, and more systematic data collection. To truly develop comprehensive urban knowledge infrastructures, municipal information systems will not simply include an inventory of existing assets, but will also include a mechanism for maintaining such an inventory.

We remain convinced that much of urban planning is an art, based on unquantifiable, instinctual, and interpersonal feelings, yet our posture should not be construed as a contradiction to our clamoring for a systematic and continuing accumulation of municipal knowledge. We also do not advocate total rational planning; we consider it utterly impossible and undesirable. However, we do believe that planning is difficult—hindered as it is by the quagmires of personalities, power-struggles, and fears—and that it would be highly beneficial to be able to rely on a solid foundation of factual knowledge as an anchored platform from which to deal with the unpredictable character of public hearings and municipal commissions. In our view, a comprehensive and well-maintained urban knowledge infrastructure is not only a useful concept, but could be a valuable reality. It is entirely feasible as long as its development is approached methodically and modularly with an emphasis on instilling an information-aware modus operandi in municipal departments. We also believe that planning professionals could play an important role in this paradigm shift that promises to make information a fundamental municipal infrastructure, on par with the pipes and cables that cities and towns have been managing and maintaining for more than a century.

## Bibliography

M. Allard, B. Amato, K. Barber, and S. Couture, *Evaluation of the Downtown Boston Parking Freeze* (Boston Project Center: WPI Interactive Qualifying Project, 2001).

M. Arun and M.T. Yap, "Singapore: The Development of an Intelligent Island and Social Dividends of Information Technology," *Urban Studies* 37:10 (2000).

I. Blizard, E. O'Connell, and E. Schmidtberg, *Building Maintenance in Quincy* (Boston Project Center: WPI Interactive Qualifying Project, 2004).

R. Brail and R. Klosterman, eds., *Planning Support Systems* (Redlands, CA: ESRI Press, 2001).

Brookings Institution, "Intelligent Middleware for Understanding Urban Markets: News from the Field," *The Urban Markets Initiative* (2006) <<http://www.brookings.edu/metro/umi/mitprogress.htm>> Accessed January 30, 2006.

Z. Budić, "Effectiveness of Geographic Information Systems in Local Planning," *Journal of the American Planning Association* 60:2 (Spring 1994) 244–263.

M. Campagna and G. Deplano, "Evaluating Geographic Information Provision within Public Administration Websites," *Environment and Planning B: Planning and Design* 31 (2004) 21–37.

S. J. Caporale, M. Cergneux, W. Freed, and A. L. Matzal, *The Development of a Prototype for an Automated Ambulance Dispatch System in Venice, Italy* (Venice Project Center: WPI Interactive Qualifying Project, 1997).

F. Carrera, "Il Traffico Acqueo nei Canali Interni di Venezia," report commissioned by UNESCO (Venice, July 1996).

F. Carrera, *City Knowledge: An Emergent Information Infrastructure for Sustainable Urban Maintenance, Management, and Planning*, Doctoral Dissertation (Massachusetts Institute of Technology, 2004).

F. Carrera and A. Hewitt, "The Premises Of 'City Knowledge: A Middle-Out Approach Toward Sustainable Municipal Data Management," *Proceedings of the 26<sup>th</sup> Urban Data Management Symposium* (Aalborg, Denmark, May 15–16, 2006).

C. Coulton, with L. Nelson and P. Tatian, "Catalog of Administrative Data Sources: For Neighborhood Indicator Systems," *National Neighborhood Indicators Partnership* (Washington, D.C.: The Urban Institute, 1997).

C. Cullen, M. Moriarty and C. Patel, *Parking Monitoring and Management in Cambridge* (Boston Project Center: WPI Interactive Qualifying Project, 2002).

J. Evans and J. Ferreira, Jr., "Sharing Spatial Information in an Imperfect World: Interactions between Technical and Organizational Issues," in H.J. Onsrud and

G. Rushton, eds., *Sharing Geographic Information* (New Brunswick: Center for Urban Policy Research, 1995).

J. Ferreira, Jr., "Information Technologies that Change Relationships between Low-Income Communities and the Public and Non-Profit Agencies That Serve Them," in D. A. Schön, B. Sanyal, and W. J. Mitchell, eds., *High Technology and Low-Income Communities* (Cambridge: MIT Press, 1998).

F. Fiorin and F. Carrera, "L'Evoluzione del Sistema dei Trasporti e del Traffico Acqueo," in Caniato et al., eds., *Venezia la Città dei Rii* (Venice: CiErre, 1999).

J. Flynn, P. Fydenkevez, B. McCarthy, and C. O'Connell, *Optimizing Parking Regulation Enforcement and Revenue Collection in the City of Cambridge* (Boston Project Center: WPI Interactive Qualifying Project, 2003).

S. Geertman and J. Stillwell, "Planning Support Systems: An Introduction," in Geertman and Stillwell, eds., *Planning Support Systems in Practice* (Berlin: Springer-Verlag, 2003).

S. Geertman and J. Stillwell, "Planning Support Systems: An Inventory of Current Practice," *Computers, Environment and Urban Systems* 28:4 (2004) 291–310.

L. Hoyt, "A Core Commitment to Service-Learning: Bridging Planning Theory and Practice," in M. C. Hardin, ed., *From the Studio to the Streets: Service-Learning in Architecture and Planning Education* (Washington, D.C.: The American Association for Higher Education, 2005).

L. Hoyt, R. Khosla, and C. Canepa, "The Power of Sticks, Leaves and Pebbles: Building a PPGIS in New Delhi, India," *Journal of Urban Technology* 12:1 (2005) 1–20.

J. Innes, "Information in Communicative Planning," *Journal of the American Planning Association* 64:1 (Winter 1998) 52–63.

J. Innes and D. Simpson, "Implementing GIS for Planning," *Journal of the American Planning Association* 59:2 (Spring 1993) 230–236.

G. Keating, P. Rich, and M. Witkowski, "Challenges for Enterprise GIS," *URISA Journal* 15:2 (2003) 23–36.

N. Kelly and K. Tuxen, "WebGIS for Monitoring 'Sudden Oak Death' in Coastal California," *Computers, Environment and Urban Systems* 27 (2003) 175–184.

R. Klosterman, "Planning Support Systems: A New Perspective on Computer-Aided Planning," in Brail and Klosterman, eds., *Planning Support Systems* (Redland, CA: ESRI Press, 2001).

S. Mandelbaum, L. Mazza, and R. W. Burchell, eds., *Explorations in Planning Theory* (Newark: Center for Urban Policy Research Press, 1996).

I. Masser and T. Wilson, "Approaches to Information Management in County Planning Authorities in England and Wales," *Urban Studies* 21 (1984).

National Research Council, *Toward a Coordinated Spatial Data Infrastructure for the Nation* (Washington, D.C.: National Academy Press, 1993).

- Z. Nedović-Budić, "Geographic Information Science Implications for Urban Regional Planning," *URISA Journal* 12:2 (2000) 81–93.
- Z. Nedović-Budić, M. Feeney, A. Rajabifard, and I. Williamson, "Are SDIs Serving the Needs of Local Planning? Case Study of Victoria, Australia and Illinois, USA," *Computers, Environment and Urban Systems* 28:4 (July 2004) 329–351.
- Z. Nedović-Budić and J. Pinto, "Understanding Interorganizational GIS Activities: A Conceptual Framework," *URISA Journal* 11:1 (1999) 53–64.
- Z. Nedović-Budić and J. K. Pinto, "Information Sharing in an Interorganizational GIS Environment," *Environment and Planning B: Planning and Design* 27 (2000) 455–474.
- F. L. Olmsted, Jr., "A City Planning Program," *Proceedings of the Fifth Conference on City Planning* (Boston, 1913).
- D. Sawicki and P. Flynn, "Neighborhood Indicators," *Journal of the American Planning Association* 62:2 (Spring, 1996) 165–183.
- J. Shuler, "Geographic Information Resource Management," *UCGIS Research Priorities* (2002) <<http://www.ucgis.org>> Accessed February 24, 2003.
- E. Talen, "Constructing Neighborhoods from the Bottom-Up: The Case for Resident-Generated GIS," *Environment and Planning B: Planning and Design* 26 (1999) 533–554.
- D. Tulloch, and J. Fuld, "Exploring County-Level Production of Framework Data: Analysis of the National Framework Data Survey," *URISA Journal* 13:2 (2001).
- E. Wilmersdorf, "Geocoded Information Incorporated into Urban Online Services: The Approach of the City of Vienna," *Computers, Environment and Urban Systems* 27 (2003) 609–621.
- A. Yeh, "Urban Planning and GIS," in P. Longley, M. Goodchild, D. Maguire, and D. Rhind, eds., *Geographical Information Systems*, second edition (New York: John Wiley and Sons, 1999).
- S. Zuboff, "Informate the Enterprise: An Agenda for the 21 Century," *National Forum* (Summer 1991) 3–7.