PART V:

"The indispensable prerequisite is then a real program.[...] If one studies the history of a town, the development of its trade and industry, and if one carefully checks all other statistical data, taking into account local needs, he will certainly have enough information to extrapolate this data into the immediate future, which is sufficient."

Camillo Sitte, 1889
"City Planning According to Artistic Principles", p. 264

The Emergence of City Knowledge

CITY KNOWLEDGE AS AN EMERGENT SYSTEM

A well-developed municipal knowledge infrastructure – as I envision it – enables city officials to not only accomplish their "normal" duties very efficiently and with less effort, through the shift from "plan-demanded" data to "plan-ready" information, but should eventually even lead to the spontaneous generation of "plan-demanding" knowledge that will inspire city officials to look into previously-unnoticed patterns or trends that may demand appropriate attention.

The needs and wants of municipalities, which have been computerizing many of their operations, and may have been dabbling with bottom-up GIS applications for the last few years, can now be met with the most recent advances in spatial management tools. Technologies such as relational databases, geographic information systems and the internet are not only affordable but also better suited than previous technologies to the new decentralized, "middle-out" approach for the spatial representation of urban features that is the focus of this paper. My own research and experience suggests that an informating approach that is neither truly top-down, nor merely bottom-up has a chance to be successful in the real world of municipal governance of urban structures and activities.

[high-resolution, fine-grained data]

Until now, the sort of fine-grained data accumulation that I propose[583] was impractical, due to the high transaction costs as compared to the perceived benefits. Today, the costs have come down and the benefits are being perceived more clearly at least by planners, who have otherwise been slow in realizing the potential of current technologies[584]. So the ratio of benefit per unit of cost is improving. Moreover, in a department-level, middle-out structure such as I propose, the primary beneficiaries of an exhaustive urban information system would be precisely those who are in charge of accumulating and updating the system's underlying datasets. The connection between cost and benefit would thus be immediately obvious within each informational jurisdiction[585].

[lower costs]

[improving benefits-to-cost ratio]

[second-order benefits]

Furthermore, the benefits are not limited to the here and now value of the mere "automation" of pen and paper processes, but actually extend into second-order, value-added realms where "informating[586]" leads to plan-demanding situations, such as the ones I presented in parts II and III.

[ripe technologies]

Current technological trends make my approach much more feasible today that it would have been as recently as the mid-nineties. In particular, the loosely-coupled nature of web-based applications, and the overall architecture of the WWW provide a natural infrastructure for my middle-out approach which is now "ripe" for real world implementations.

---

583 And I have applied both in Italy and in the U.S., as shown in Parts II and III.

584 Geertman and Stillwell, 2003 and 2004.

585 Singh, 2004.

586 The "automating" vs. "informating" dualism comes from Zuboff (1991). These more sophisticated benefits will require coordination to enable cross-department sharing, which would add some cost and complexity.

the emergent nature of City Knowledge

City Knowledge systems are "complex adaptive systems that display emergent behavior"[587]. Just like the use of local knowledge by individual ants informs microbehaviors that create the macrointelligence of the ant colonies, so too bottom-up, department-level data and information collected for dealing with specific aspects of municipal affairs can produce city knowledge that supports higher-order urban analyses and decisions. It is at this higher level that planners are meant to operate, instead of being forced to hunt down the data they need over and over again. Middle-out is my proposed pathway to get us there from here.

So, with costs greatly abated, and demonstrable benefits that reliably exceed the first-order advantages of computerization, I conclude that now is a good time to begin to treat information as a *bona fide* infrastructure and I therefore recommend to progressively create a City Knowledge system from the finest grain out.

analyzing emergence

Even though I have not fully explored the problems and promises of the extension of City Knowledge beyond a single department, many of the department-level information systems that I was involved with – especially in Venice – are now mature enough to be evaluated in such context. Such an evaluation would be a useful follow-up to this dissertation since it would allow us to: identify the practical impediments to the horizontal (and vertical) diffusion of City Knowledge across department boundaries; to better quantify the costs associated with the added coordination necessary to cross-reference datasets from different organizations and to keep them synchronized; and to address some of the transition issues facing municipal departments when they try to mainstream promising technologies into established work flows and entrenched bureaucratic processes. For the time being, I will conjecture what

creating emergence

In this final Part V, I suggest that instead of "dealing with" emergencies, city officials may actually purposely set up an *emergence* of their own. Instead of waiting for the arrival of a magic, all-encompassing system for the management of all municipal information, I recommend a path that promises a non-traumatic transition to City Knowledge in true *emergent* fashion. I hypothesize how cities can start from one specific aspect of municipal affairs and then sustain such a system in perpetuity, gradually spreading this approach to more and more aspects, then infecting other departments and eventually diffusing into other towns, so that this concept will become commonplace in urban communities around the world.

[587] Johnson, 2001.

CREATING CITY KNOWLEDGE FROM THE MIDDLE OUT

Geographical Information Systems (GIS) and Database Management Systems (DBMS) – while rather commonplace in today's city government – are scarcely used to systematically keep track of essential urban elements, such as roads, trees, sewers, etc. in an "automatic" fashion. Ad hoc, or "implicit" knowledge is used instead to produce maps and datasets that feed the decision-making process on a need-to basis, case by case. The typical city department is generally too busy with the struggles of day-to-day emergencies to afford the "luxury" of thinking ahead and planning a systematic knowledge acquisition campaign that will last several years.

Yet, this study recommends that cities implement just such a long-range infrastructural approach to city knowledge. How can one realistically expect to accomplish such a demanding – albeit promising – task, when the average city official is burdened with myriad challenges just to perform routine activities that the administration of the city demands? This dissertation proposes to towns to embrace the piecemeal, bottom-up, opportunistic evolution of municipal information systems that results from our "muddling through", but to do so strategically and deliberately, with an eye to the overall role of each department's information in the combined big picture of City Knowledge.

In this section, we take a look at a plausible implementation path that cities and towns could adopt to harness urban data used for day-to-day maintenance into plan-ready information that can facilitate more complex management tasks and can lead to the kind of deep knowledge that can inform longer range plans for the future of our communities.

The basic steps that an office or department could follow iteratively to move toward a City Knowledge system are:

(1) commit to a City Knowledge approach to the development of a comprehensive municipal information system and accept the middle-out strategy as a possible tool to achieve it;
(2) define its informational jurisdictions by analyzing information flows;
(3) identify "low-hanging fruits" within its jurisdiction – i.e. applications that would yield the maximum benefits at the lowest cost;
(4) collect "atomic" data about the urban elements related to the low-hanging fruits;
(5) develop and apply a mechanism for sustainable data updates;
(6) share the information and data as appropriate;
(7) coordinate as needed.

In the sections that follow, we first discuss the need for a commitment to the development of City Knowledge, which can be manifested by an official acknowledgement of the infrastructural role of information. Being fully aware that this step is the most important and least trivial, we subsequently assume that this paradigm switch has taken place in

order to discuss all of the rest of our strategy and tactics. The hope is that, in accordance with the theory of communicative action[588] and the "tipping point" theory[589], this dissertation and possible future pilot implementations of City Knowledge systems will make it more and more spontaneous for municipal offices to make such a commitment, which for now we will only assume.

Of course, as is often the case, the devil is and will continue to be in the details. Just where one draws the line when it comes to key decisions regarding the application of these principles will depend on the circumstances of that particular department at that particular moment. There is no one-size-fits-all silver bullet when it comes to City Knowledge. As long as there is a commitment, though, I think progress can be made by any town anywhere, with only a modicum of computer savvy and human and financial resources.

The rest of this chapter fleshes out a realistic pathway to City Knowledge using the "lessons" from parts II and III as supporting evidence and as examples of not only what worked, but also of the thornier implementation issues encountered or still to face.

The next and final chapter provides insights in how to keep these systems alive and how to make them become more common across the municipal landscape. The next section begins to address the question of exactly how to construct such comprehensive municipal knowledge.

Commit to City Knowledge

A desire to create a comprehensive municipal information system has been around for decades[590], so for planners and municipal practitioners the concept of City Knowledge is not new. My thesis takes for granted that towns will want to use a comprehensive municipal information system.

What I think is new in this dissertation is the approach I propose for how to get to the holy grail of full city knowledge gradually and methodically as long as the town has made the simple commitment to become "information-conscious".

Some cities[591] have attempted to push a wholesale approach from the top down, forcing departments to comply with directives from on high. Almost everywhere, in the meantime, we have witnessed the spontaneous appearance of homegrown GIS endeavors at every branch of the municipal hierarchy in medium to large cities, primarily but not only in the "developed" world. I would like to contend that so-called "top-down" approaches have failed in their ability to percolate downwards just as bottom-up initiatives have failed to bubble upwards. In essence, it is rare (if not impossible) to find comprehensive GIS infrastructures within municipalities that are coherent and interconnected at all levels. If these two popular methods for

top-down

bottom-up

588 Habermas, 1984, 1987; Innes, 1995.

589 Gladwell, 2000.

590 Geertman and Stillwell, 2004; Brail and Klosterman, 2001.

591 Like Singapore (Arun and Yap, 2000) and Vienna (Wilmersdorf, 2003), but also my own hometown of Venice (with the SITEA system).

the creation of a knowledge infrastructure have largely failed, what alternative is there?

middle-out

As amply discussed in a previous chapter[592], I propose a hybrid approach to knowledge acquisition and maintenance that combines the freedom, flexibility and creativity of the bottom-up scheme with the structure, rigor and standardization of a top-down system. I call this a "middle-out" approach. This method engages offices or departments at the "front-line" of city operations, but avoids the pitfalls of the extreme, independent and individualistic idiosyncrasies associated with random homespun initiatives. In a middle-out approach, there is coherence and consistency within the front-line office; there is coordination between departments; and there is a conscious appreciation of the broader context in which information is being managed above and beyond the immediately obvious needs of a single office.

A committed city will treat information as an infrastructure and that ought to make all the difference in the world in how information is regarded in that town from that point on. Such a commitment would provide a strategic thrust for the gradual accrual of City Knowledge, which will be aided and abetted by the tactical lessons included in this chapter and throughout parts II and III.

Just what would make a town commit to this paradigm shift remains a question for further study, though the benefits derived by small-scale applications may incrementally build up the case for City Knowledge and could eventually bring about this sort of commitment from small towns as well as large cities, from Venice to Vancouver, from Peoria to Peking. Every town will have its own threshold beyond which the presumed costs are offset by the perceived benefits of such a commitment. Each department in each town, in fact – according to my middle-out strategy – will be facing this quandary and will have to "draw the line" between committing to this approach or not.

Identify jurisdictions

According to our scheme, municipal information ought to be divided up among all functional divisions or departments in the city, along informational jurisdictions. We propose to "deconstruct" the urban landscape into homogeneous slices and "atomize" the data collection within each jurisdiction to achieve the smallest grain compatible with realistic needs and available financial and human resources.

As explained in detail in a previous chapter, "birthrights" and "deathrights" could be identified by analyzing information flows in typical departmental procedures[593]. It is preferable to pick an informational jurisdiction that is as unequivocal and uncontroversial as possible for the first City Knowledge application. You don't want to start by stepping on someone's toes (or turf) right off the bat.

592 Page 196.

593 As was done by Hart *et al.*, 2004 and Novello and Sartori, 2004.

This is the most normative of my tenets since I am convinced that this step is the most important for the long-term prospects of any municipal information system to flourish and persevere, as our *lessons from underground data management* [594] and many other positive[595] and negative[596] experiences have hinted at. This aspect will require some coordination among agencies to define the domains and "draw the lines" of demarcation between elements of the urban environment that are under each department's jurisdiction.

There will undoubtedly be turf battles over these choices and some urban elements or activities may even remain unassigned to any department, either because nobody wants them or because two or more offices lay claim to them. My view is that the latter controversies can be easily solved by letting either department – or even both – take control under strict performance expectations, with clear, yet reasonable deadlines for the delivery of shared datasets and layers that the "other" overlapping department can immediately share and use. I would let the department that lost the jurisdiction battle be the judge of the adequacy of the portion of City Knowledge that the "winner" was able to put together by the agreed deadline. If the product was deemed inadequate, then I would give the dissatisfied department a chance of doing a better job, under similar deadlines. This tension may actually produce much better results than unopposed assignments of domains.

If nobody wants to take charge of a particular aspect of the urban landscape[597], then some office may be put in charge of it *pro tempore* until the matter can be resolved by either determining that such an aspect is not worthy of special attention, or by identifying institutional reasons to attribute the onus to a specific office. If anybody ever comes looking for data about this undesirable domain, then they should be either be put in charge of it – if appropriate – or they should be probed in order to arrive at a resolution of the issue. If nobody seems to use information about the elements in this category, then there is no need to worry about this domain until somebody does.

Future research on this aspect may provide more details about the most appropriate consensus-building approach to the formation of these important jurisdictions.

---

[594] See page 103.

[595] All of our canal work fell clearly under Insula's jurisdiction, which greatly facilitate our efforts at making the data usable and re-usable. See part II.

[596] The unclear boundary between the *Soprintendenza ai Beni Artistici e Storici* and the *Soprintendenza ai Beni Ambientali e Architettonici*, as well as the even more controversial private vs. public jurisdiction over public art has been the primary cause of the neglect suffered by the over 4,000 pieces of art visible from the streets of Venice. See page 112 for more on Public Art.

[597] As was the case in Venice with Public Art.

Pick Low-Hanging Fruits

Having identified the informational jurisdictions, we recommend that each department pick a specific process within it, and target the urban structures or activities that participate in that process. The selection of the "low-hanging fruit" should be based on benefit vs. cost[598], as well as on the administrative priorities of the particular political moment. Ideally, the first project should be a "no-brainer"[599], i.e. one that actually pays for itself[600] or even makes the town money in the short term[601]. The choice of "grain" ought also to be tailored to the needs of the moment so as to make the project a 'showcase application".

Based on the lessons we learnt in Venice and in Massachusetts, each department should have at least one departmental body of information that it is already compiling, and has been for years[602]. Perhaps that ought to be the first candidate for a full-cycle City Knowledge review to see how to set up a system that would keep track of those data in perpetuity, including all of the periodic changes that might occur in the future. When dealing with cultural, historic assets[603], we have also demonstrated the benefits of acquiring plan-ready data for finite datasets of elements that are permanent, or very slowly changing. These datasets may be the lowest-hanging fruit and provide good returns on investment since they need to only be collected once and the cost of updates in minimal.

Once again, it is non-trivial to decide exactly where one draws the line between what's low-hanging and what isn't. Such decision would most likely be based on some cost-benefit calculus, which itself requires a baseline of information that may or may not be available. When it comes to inventories of municipal assets, it is always difficult to determine the cost of the complete catalog, since cities usually do not know how many items that inventory actually contains. Therefore, contracts such as those for our work on bridges and docks in Venice[604] are based on a per-item charge[605] and the city every year could just budget an amount that corresponds to *n* objects in the inventory, and so on until the job is complete.

More research could be done on the meaning of "return on investment" (ROI) in this context. The research would have to match perceptions about costs and benefits to choices with an eye to the process of selection and the definition of "low-hanging".

---

598 See the ample discussion of this important step in Reeve and Petch, 1999, pp. 51-81.

599 The "showcase application" mentioned in Reeve and Petch, 1999, p. 57.

600 As our Cambridge parking project (Flynn *et al.*, 2003).

601 As did our Parking Freeze fee proposal (Allard *et al.*, 2001).

602 See parts II and III.

603 See [lessons from palaces, churches and convents] on page 109 and ff.

604 See part II.

605 As you may recall (page 169) the rule-of-thumb unit cost for complex urban assets such as dock and bridges came out to about $150 per element to collect and archive attribute data as well as GIS attributes.

Atomize and Conquer

Once the department has selected the specific municipal process to "informate", the next task is to dissect the process to arrive at the "atomic" elements about which to collect information in a systematic and perpetual way. Here is yet another situation that forces a decision on where to draw the line. Whatever our "atom" is today, it may be too big to represent tomorrow's sub-atomic particle. How fine of a grain should we use? What should be the unit of measurement? The answer, once again, is not easy[606].

atomic grain size

First of all, there are two grains to consider here: the size of the spatial units that will be managed and the detail collected for each of the spatial units. My quick choice would be: smallest area and largest attribute set. But, once again, you need to draw the line somewhere.

[spatial grain]

For the spatial dimensions, the expected unit of analysis should be one factor guiding the choice of unit of measurement – which should always be smaller than, or equal to the former. Cost considerations will dictate the grain to some degree, but the choice will be mostly practical and commonsensical, based on the physical object itself, as we learnt in Venice[607].

[attribute grain]

The choice of detail in the attributes to collect for each unit is less straightforward. As our cases showed repeatedly, there are advantages to be reaped by collecting as fine a grain as time allows[608]. This enables re-usability since finer grain can recombine more flexibly. Moreover, the attributes ought to be true to reality and not synthetic expressions of expertise, as we encountered with our tree management and public art projects. Instead of mentally calculating an overall condition rating for a tree (good, bad, OK), we collected visual clues of each factor that might play a role in the overall tree condition, such as insects, fungi, root problems, structural problems, etc. This adds many attributes, but makes the work more manageable and probably cheaper in the long run[609]. Beyond the separation of facts and values[610] in the choice of attributes to characterize the elements being inventoried, there is one last consideration one ought to make before the data structures are ready to be populated with field work or otherwise.

[attribute enrichment]

I am referring here to the teleological enrichment of the dataset through the choice of additional parameters that may not be needed right away for the task at hand, but are not too hard to collect once the crew is standing in front of the object in the field, and they may have some future use, either within the office itself or for some other maintenance, management or planning activity. This is where the science turns to an art. The line one draws here is very blurry and its shape can change dramatically depending on who is drawing it. A choice of additional fields made by someone with a planning mindset may permit a more artistic palette of re-uses and more sophisticated second-order capabilities, whereas the choice

606 See the lessons we learnt when choosing between using buildings ort floors as atoms of City Knowledge in Worcester (on page 156).

607 See [Lesson 1: atomize] on page 50, and other [lessons from our physical studies] on page 62.

608 See [re-usability of fine-grained data], [use finest grain], and [keep data in raw form] on page 68 ff.

609 See page 126 and ff.

610 See the lesson entitled [separating facts and values] on page 131.

made by someone with a management mindset might be more scientific and perhaps less creative and less combinatorial, but more concrete and focused on specific foreseeable future needs and wants.

locus of control

Plan-ready information will be generated more efficiently if the work of creating and maintaining the databases, GIS layers and information systems is left to the "frontline" offices of each department, where the real action is and where municipal data are acquired and generated. Frontline offices in turn should generally attempt to delegate most data collection and updates to outside contractors or to the public and, only as a last resort, they should devote municipal staff to the task if necessary. Since the data would be collected first and foremost to suit the needs of the front-line department, their collection should be verified by field inspections by town officials to guarantee the integrity and quality of the data.

[collect fine-grained data within jurisdictions]

[sustainably maintain data at frontline]

[adopt standards and methods]

Once enough data have been collected (using agreed standards and methods), and organized into plan-ready information, sharing with other agencies may be possible either by operating along existing interagency lines of cooperation or by crafting new virtuous networks of value-added knowledge.

[share to leverage value-added benefits]

data collection

Having settled on the spatial and attribute grain, and having identified the frontline offices in charge of the data collection, as well as the surrogates who will inexpensively collect the data for those offices, a town that has embraced the cause of City Knowledge faces two fundamental tasks:

- Catching up with the backlog of urban data
- Dealing with future changes in the city

catching-up with the backlog

In my middle-out approach, departments (or even divisions or offices within a department) are responsible for their own information. The fact that the department or unit takes ownership over its own knowledge guarantees that the data are *reliable, complete, pertinent* and *up-to-date*, since the job of department members rides on the quality of these data. The department needs to make a commitment to the acquisition and upkeep of its slice of City Knowledge before any of the following steps can even be envisioned.

[commitment to city knowledge]

paying for the backlog

It is undeniable that an initial collection of data will be necessary to catch up with the backlog of information that is already "out there" in the city today. Naturally, towns will explore the no-cost or low-cost solutions first.

[no-cost solutions]

The no-cost solutions entail getting data from third-parties for free. No-cost is probably an illusory concept since the mere acquisition of pre-existing data would itself entail some cost, even if the town were to receive the "complete municipal information system" as a "free-gift"[611].

[free = self-reporting]

Data could be collected more or less for free by mandating a lot of self-reporting from the public and from businesses. This may include

[free = required information returns]

[611] In this section, I will call no-cost or free what is actually a fairly low cost activity, but certainly not entirely free-of-charge.

making information-returns part of the current requirements for existing processes (like permit applications, for instance). Similarly, contractors or professionals hired by developers submitting a plan, or hired by the town[612] to rebut the same plan, could be forced to submit information in a certain format in order to feed the growing information infrastructure[613].

[free = secondary sources]

Free data could also come from secondary sources, be they government agencies (like the regional planning agencies, or the Census Bureau) or private or semi-private entities (like Chambers of Commerce, Real Estate professionals organizations, etc.). Acquiring these data will necessitate some formal agreement but government-to-government data exchanges ought to be free and fairly feasible with a modicum of bureaucracy and legal paperwork. Post-processing of these datasets may involve getting rid of "stubborn errors", but the guaranteed reliability of the data may justify investing in middleware such as Ferreira (1998) proposes[614].

[free = volunteers or students]

A final source of "free" primary (or secondary) data can be students, researchers or volunteers. All of my examples come from the "free" work of WPI students and even from volunteers (*Earthwatch*, grade schoolers, interns). Data validation and quality control are factors to be considered carefully here[615]. Involving local citizens in these efforts produces a multiplier effect.

[low cost= maintenance-based data collection]

Low-cost, very gradual ways of slowly chipping away at the "backlog" may include collecting data as repairs or maintenance are done. If the repair work is farmed out to contractors, a specific informational return would be included in the bid language and in the subsequent contract.

[low cost = contractual obligations]

Another way for a town to gather information gradually – yet unsystematically – is to force its contractors to produce GIS-compatible data whenever some spatially-referenceable service is rendered[616]. Towns could simply change the wording of contracts and bids to impose informational deliverables upon external contractors[617]. However, before doing so, that town needs to make sure it is ready and able to deal with such an information onslaught once it has mandated it.

[low cost = in-house]

A final recourse would be for the department to do some of the data collection with its own staff. This is not a preferred path, but could be considered, depending on the staffing levels and on how busy municipal staffers are. Some institutionalization of the process of catching up with the backlog may be instituted by including data-gathering as one of the requirements in the job descriptions of municipal workers.

---

612 Often at the developer's expense.

613 As we did in Gage *et al.*, 2003 and in Brown and Groeli, 2003. These information-conscious requirements may provide a net gain for the town as long as these costs are not somehow transferred to the community in some other way, like through higher real estate costs.

614 Ferreira, 1998.

615 For example, Forma Urbis ended up re-collecting all of the bridge data for Insula, despite the fact that half the bridges had already been measured by Bahn *et al.* in 1998. Generally, Forma Urbis will always collect the data again if possible. The "methodology" is therefore the primary contribution that student projects really make to the final operative information systems, in my experience.

616 For example, surveyors may be asked to do that, as well as architects, or professional planners.

617 This would not be free because it is presumed that the cost estimates would rather quickly incorporate the informating costs and pass them onto the town.

[market cost = consultants]

Catching up with the backlog may, in some municipal domains, may be urgent enough or sensitive enough to warrant farming out to outside consultants to expedite matters along and get the program off to a good start so it can act as a model and a catalyst. Capturing "snapshots" of *activities*, such as traffic or housing patterns, may also entail the participation of professional consultants.

[information budget]

Regardless of the route a city decides to take, the town will incur some costs. These could be offset, at least in part, by devoting a fixed percentage of a department's regular budget (say 5%) to these knowledge-acquisition activities, thus making the process rather *affordable* during any FY.

Future research could focus on some of the more creative no- or low-cost data-gathering schemes mentioned herein, in which we have little or no direct experience and are not really discussed in the literature either. In particular, I plan to experiment more cogently with the *quid pro quo* bartering of information[618] and the tweaking of contracts and job descriptions in Venice and in Massachusetts.

[618] Between Spencer and the Central Massachusetts Regional Planning Commission. See footnote number 437.

Update Sustainably

All municipal activities leave a paper trail that would easily lend itself to automation. Roads are regularly re-paved, cleaned and cleared of snow, so someone is issuing work-orders or stipulating contracts for these services. Similarly, trees are bought, planted, removed and trimmed and paperwork is produced to make each of these actions happen and to keep track of the corresponding expenditures. Sewers, like many other components of the urban infrastructure, are subject to similar record-keeping procedures, plus they are also regulated and licensed. The list goes on.

automating and informating

Not only should these informational opportunities be tapped into, as much as possible, to populate the city's knowledgebase as we catch up with the backlog, but they should play an even more important, primary role in keeping the information up-to-date to promote the sort of "automation" that would lead to a more rigorous approach to informating urban maintenance, management and planning[619].

Just about all of the techniques listed in the previous section that could save a town some money in the collection of data, could also apply when it comes to data updates. The most important and most effective approach to sustainable information upkeep would be the institution of computerized mechanisms to intercept changes in the urban realm as they happen, never creating a backlog ever again. This step would entail considerable thought and planning, together with technical, legal and administrative agreements and solutions.

[capturing transactions]

[update or die]

The advice to a community starting on the path to City Knowledge is to never embark in a system to keep track of a class of assets unless there is a plan to capture future updates too. Updates should be sustainable because they can be more expensive than the backlog in the long run, and they can make or break a project down the line. Some projects are born obsolete because of their disregard for the upkeep of the datasets.

[experiment with innovative updates]

The sticky points as well as the interesting research angles are about the same here as they were for the collection of the backlogged data discussed in the previous section. The next step for me in this context will be to go back to the more mature applications that I had a hand in starting in Venice and in Massachusetts and take a closer look at gaps in the updating mechanisms, so that we can experiment with some of the novel techniques that are introduced in this dissertation but are still mere conjectures and hypotheses until we put them to the test in the real world.

619 See the lesson about [maintenance-based updates and cataloging] on page 131.

Share Appropriately

The full power of City Knowledge and its multiplicative, emergent effects will not become tangible unless some amount of sharing takes place, to bring more than one department, agency or institution into the system.

[share to foster emergence]

As soon as a department has mastered a specific application for the maintenance, management or planning of some urban aspect over which it has jurisdiction, the information is available for sharing if so desired. This step is not necessary, nor mandatory, but it would be certainly salutary, so I would personally encourage the practice of sharing non-controversial, privacy-respecting data, in order to generate the amalgam within which a fully emergent City Knowledge system can flourish and produce unexpected, value-added benefits[620].

The technology itself, once again, is not the solution, nor the problem. But it may contribute to a more rapid transformation of the internal organizational dynamics of government agencies toward a "connected distributed" [621] *modus operandi*, that will enfranchise the citizenship as well as the front-line civil servants. This, in turn, may set the stage for a truly devolved informating "wholeness"[622], where managers and managed contribute together, "holistically"[623], to a middle-out approach for the management of urban affairs. The Internet and the WWW would certainly facilitate such a development.

All of these innovative approaches may incrementally lead to "*a truly interactive, timely planning dialogue between neighborhood planners and city agencies – as well as* [to] *a mode of interagency* [and – I would add – 'intra-agency'] *coordination that might allow agencies to keep pace with one another*"[624] and with their public constituency. Once again, the interconnectivity provided by the WWW today makes this interactive approach all the more feasible and affordable for our cash-strapped public agencies.

[share to avoid duplication]

Another – more practical – reason for sharing is to avoid duplication of efforts within the department or the municipality[625]. Of course, the avoidance of duplication implies that some method exists to inform all potential users of the existence of urban data and information for such and such an asset or activity over such a span of time and over such a spatial extent. This is what metadata is really good for. Unfortunately, the use of metadata – even in progressive GIS-intensive municipalities – remains lackadaisical so far, especially at the real front lines. Even my groups have not yet used any part of the FGDC metadata standard rigorously or methodically.

Metadata use and usefulness is an area of research that certainly deserves more attention and I plan to look at my examples through this lens as well, when I go back for the aforementioned evaluation of mature systems.

---

620 As was the case in the [lessons from palaces, churches and convents] on page 109.

621 Thomas W. Malone, *Is Empowerment Just a Fad?* (1997)

622 Zuboff's term (1991).

623 This term borrowed from Evans and Ferreira, 1995, p. 458.

624 Ferreira, 1998.

625 See the lesson entitled [share results to avoid duplication] on page 109.

Coordinate as Needed

Once enough sharing takes place – most likely along paths of least resistance connecting two agencies with pre-existing administrative ties – the issue of coordination will immediately come up. While this is not a mandatory step, it is very likely to be required once the number of sharing agencies becomes greater than two. In fact, many issues will become thornier once interoperability is more in demand across departments.

Coordination is costly and needs to address dependencies and problems with synchronization. Sharing also entails dealing with many other subtler intricacies such as the issues of data quality, reliability, liability, copyrights, accessibility, security and many others. All of these issues require coordination.

[grow into coordination]

I am purposely downplaying the importance of coordination because I don't want to scare people from engaging in all of the other steps listed before. In fact, I was not going to even include this step at all for that reason. Many people react negatively to too much oversight and too many rules – including me. It would be great to establish coordination committees or roundtables within a municipality, but I think that the membership into these coordinating entities should emerge slowly (like everything else in my City Knowledge approach).

[coordinate spontaneously]

I believe that individual departments need to grow their internal information capacity and begin two-way sharing until they themselves experience the self-generated impulse to coordinate, at which point they can be ready to join the interdepartmental coordinating club. Before that time, they may participate in meetings as observers to become acquainted with the issues, but the methods[626] will have no "stickiness"[627] until the urge to coordinate emerges spontaneously from within the department.

There are many ripe areas of research in this arena, but I am not interested in the technical standards per se, but in how to make them useful and used. I would like to apply some lessons from the field of emergence[628] to determine what minimal sets of standards and protocols are needed to bring out the second-order, high-level emergent qualities of City Knowledge, without overburdening the users.

626 Described starting on page 219.

627 Gladwell, 2000.

628 Johnson, 2000. I am especially intrigued by T*ipping Point* (Gladwell, 2000) and *Ideavirus* (Godin, 2001).

A SUSTAINABLE FUTURE FOR CITY KNOWLEDGE

Despite all of the great cases presented herein, and despite the continued success of our initiatives in Venice and in Massachusetts, to date none of these municipal administrations has wholeheartedly embraced City Knowledge as its official *information-aware modus operandi*. There have been comforting signals from both sides of the Atlantic[629], yet there is no proof out there that these concepts would fully work outside of the privileged position that I have put myself into – namely of having the luxury of coming at problems laterally, with considerable human and technological resources and without needing to be accountable to bosses, citizens or taxpayers.

One of the most important future developments – hopefully in the short run – will be to have some real examples of cities and towns (or at least departments) that formally espouse these principles and begin to deal with the devilish details of implementing and keeping alive such an operation. There will be innumerable lessons to learn once these experiences start in full earnest.

the tipping point of City Knowledge

Since it is an emergent system, I think that the small seeds of City Knowledge that are already germinating in Venice, Cambridge, Boston and Worcester will gradually grow to encompass entire departments and then infect the whole city. As Malcom Gladwell explains in *The Tipping Point*[630], after these innovators and "early adopters" have established the concepts in their respective organizations and communities, we'll need to wait for the "connectors", "mavens" and "salesmen" to do their subtle work before we can witness the emergence of a more sizeable "early majority". This "tipping point" will occur when the adoption of City Knowledge reaches the "moment of critical mass, the threshold […] where the unexpected becomes expected, where radical change is more than a possibility". When or whether this will happen depends on the communicative skills of practitioners like me who can see the possibilities that plan-ready information holds for planners and for decision-makers and can articulate a strategy to achieve it[631].

value-added benefits for high-order users

Planners can reap the value-added benefits of being able to jump right into second-order analytical tasks – which is what they are probably best at, in addition to being what they are paid for – without having to track down all of the necessary data first. Yet these higher order capabilities will only become available once the front line municipal workers begin to see direct first-order benefits for themselves and their job. It is the day-to-day value derived from the automation of the front-office activities that enables the back-office to harvest the deeper informating benefits. I posit that this useful commingling of instant gratification with long-term capacity building

value for front-line users

---

629 The latest was the honor of having been invited to talk about City Knowledge principles to the top echelon of the M.I.S. department of the City of Boston (on 8/7/04). Meanwhile, in Venice, Insula continues to carry the torch as the active operation that most closely resembles City Knowledge at work.

630 Gladwell, 2000.

631 For example, treating City Knowledge as an *ideavirus* (Godin, 2001) may yield some useful insights.

middle out approach

can be achieved with our middle-out method that mixes the best features of bottom-up and top-down to create a sensible, low-impact approach to the accrual and sustainable maintenance of city knowledge.

cost vs. benefit

In my view, a City Knowledge system is both desirable[632] and feasible[633]. My personal cases document how the emergent power of plan-demanding knowledge provides informating benefits that are higher than the typical improvements due to mere automation. The technological costs of such an endeavor are declining every day, whereas the organizational, administrative and logistical costs can be kept under control, using the techniques that many of my cases have demonstrated in the real world. In short the balance between benefit and cost is tipping toward the former.

toward city knowledge

These favorable circumstances make finer-grain data collection conceivable and affordable, which in turn should yield more flexibility in the re-utilization, aggregation, manipulation and analysis of our urban datasets. With only a modicum of overarching coordination, different departments can set out to comprehensively capture the data that they need in order to maintain assets, manage activities and plan developments within their informational jurisdictions. As they catch up with the backlog of information already out there – starting from low-hanging fruits[634] – they can also begin to put in place self-perpetuating mechanisms for the upkeep of the fundamental framework datasets that are necessary to fulfill their own departmental needs. The resulting plan-ready systems promise to make each office perform its duties more effectively, efficiently and efficaciously.

Following the path of least resistance, these plan-ready systems can then be shared where institutionally or opportunistically appropriate. I hypothesize that sharing can give rise to the self-organizing behavior that City Knowledge has already demonstrated, allowing plan-demanding situations to emerge as we gain a deeper and deeper understanding of the urban fabric in which we live and operate.

the future as a possibility

With this dissertation, as a reflective practitioner, I have put forth my new metaphors for the redefinition of City Knowledge as a paradigmatic state-of-mind. I hope this concept will penetrate into the collective subconscious of municipal governments, so that they will begin to treat urban data, information and knowledge as essential infrastructural elements of our cities and towns.

As a communicative planner, I hope that having defined, proposed and partially demonstrated a framework for the gradual growth of City Knowledge, it will no longer be a vague holy grail that we could perhaps achieve "some day", but will instead become a clearer cause to champion for those of us who believe in its power for the transformation of communities.

---

632 As indicated in Marchewka, 2003, pp. 8-9.

633 As my cases in Part II and III showed.

634 i.e. the element of the urban realm that represents the best return on investment weighing the transaction costs versus the timeliness and value of the benefits.

In essence I am proposing to replace the tactical process of "muddling through" with a strategic equivalent: the conscious and deliberate "middling through". I think technology will fuel this process. For example, I foresee a future in which mobile technology will not be focused exclusively on "Push"[635]. I am sure that we will soon witness the emergence of "Pull", with which we will tap into the data-collection potential of these mobile devices[636]. I think towns ought to be the first to "pull" geo-spatial data from the field using these cutting edge technologies at the very front of the front lines.

In addition to technological developments, there are many other areas of research that promise to be useful to give more body to the fledgling concepts of City Knowledge and middle-out here introduced. I look forward to deepening my understanding of these mechanisms, so that towns can really find smooth ways to acquire the data, information and knowledge that they require for urban maintenance, management and planning.

As I said in the introduction, I hope that some day City Knowledge will itself become transformed from a plan-demanded cause pushing us from the past toward a future full of emergencies, into a different type of cause – a plan-demanding cause celeb – pulling us into a future full of emergent possibilities for our cities and towns.

Ten or twenty years from now City Knowledge and middle-out may or may not be household words in municipal administrations, however, if cities and towns will be making progress towards comprehensive municipal information systems that have the qualities I listed in this dissertation, then perhaps the paradigm shift will have taken place – perhaps under another name or with a different approach – and, if things go the way I predict, these concepts may be so ingrained in the municipal psyche to be virtually unnoticeable. I hope to be around when that begins to happen.

635 Whereby connected users are able to receive information about everything everywhere (from stocks to location-aware advertising, news, media, etc.) through wireless PDAs and tablet PCs. Push was on the cover of Wired magazine way back in March of 1997, but there has recently been a resurgence of interest in it.

636 I am advising a project on "Pull" in the 2004-2005 academic year. As far as I know this may be my own neologism. I haven't heard it used in juxtaposition with Push.