PART IV:

"The enormous importance of such a City Plan Office as we have been discussing, with its elaborate, active and obviously costly human machinery for systematically recording these live ideas which form the real city plan, for interpreting them and for deliberately amending them, lies in the fact that without such machinery these functions are performed unsystematically, intermittently and very imperfectly by people whose principal interests and duties lie in other directions. Without it the actual set of ideas and purposes concerning probable future improvements and conditions which are really kept in mind in such a way as to have practical influence upon current decisions, is dependent upon the memory and personal equation of scores of different individuals, no one of whom has opportunities to be cognizant of the whole field or to keep in touch with all the other people."

Frederick Law Olmsted Jr., 1913
Proceedings of the
Fifth National Conference on City Planning, Boston
http://www.library.cornell.edu/Reps/DOCS/olmst_13.htm

THE ESSENCE OF CITY KNOWLEDGE

A TALE OF TWO CITIES: STRUCTURES AND ACTIVITIES

" The city of Sophronia is made up of two half-cities. In one there is the great roller coaster with its steep bumps, the carousel with its chain spokes, the Ferris wheel of spinning cages, the death ride with crouching motorcyclists, the big top with the clump of trapezes hanging in the middle. The other half-city is of stone and marble and cement, with the bank, the factories, the palaces, the slaughterhouse, the school, and all the rest. One of the half-cities is permanent, the other is temporary, and when the period of its sojourn is over, they uproot it, dismantle it, and take it off, transplanting it to the vacant lots of another half-city.. And so every year the day comes when the workmen remove the marble pediments, lower the stone walls, the cement pylons, take down the Ministry, the monument, the docks, the petroleum refinery, the hospital, load them on trailers, to follow from stand to stand their annual itinerary. Here remains the half-Sophronia of the shooting galleries and the carousels, the shout suspended from the cart of the headlong roller coaster, and it begins to count the months, the days it must wait before the caravan returns and a complete life can begin again...

Italo Calvino
"Invisible Cities", p. 63

Kevin Lynch – a professor in the same MIT department where this dissertation was developed – who remains, despite his premature death, a role model for those of us who are interested in the phenomenological aspects of a city in relation to planning, once said:

> *"The Fundamental problem is to decide what the form of a human settlement consists of: solely the inert physical things? Or the living organisms too? The actions people engage in? The social structure? The economic system? The ecological system? The control of the space and its meaning? The way it presents itself to the senses? Its daily and seasonal rhythms? Its secular changes?*
> *Like any important phenomenon, the city extends out into every other phenomenon, and the choice of where to make the cut is not an easy one".*[341]

His resolution of this dilemma was simple and straightforward. As he put it:

> "[…] *the chosen ground is the spatiotemporal distribution of human actions and the physical things which are the context of those actions* […]".[342]

My shorthand way of rephrasing Kevin Lynch's dichotomy is simply that cities are made up of two components: *structures* and *activities*[343]. Italo Calvino eloquently captured the complementarity of permanent structures and ephemeral activities by splitting in half the city of *Sophronia*, my favorite of his *invisible cities* (left).

structures

Structures include all of the "containers", environments, spaces and places that make up the physical, material city. These concrete components of the urban realm include buildings, parks, rivers, roads, trees, fire hydrants and everything else that's "out there" in our cityscapes. Structures lend themselves to "permanent" inventories since they change ever so slowly and can thus be captured once and for all through an initial cataloguing effort, only to be occasionally updated by intercepting administrative acts that signal the changes that do occur – however seldom – in the physical make-up of the city.

activities

My dissertation deals predominantly with the physical structural elements of municipalities, though *activities* are also discussed at length. These dynamic phenomena are more difficult to track because of their ever-changing nature[344], which makes it impossible to capture the information

341 Lynch, *Good City Form*, p. 48.
342 *Idem.*
343 Hopkins, 1999, p. 335 offers essentially the same breakdown in Figure 1.
344 See however Longley and Harris, 1999.

once and for all, and then only deal with small changes from that moment on – as is instead the case with our urban structures. Whereas information about *structures* can be maintained through a piecemeal "transactional" approach, which is less costly and more manageable on a day-to-day basis, *activities* need to be monitored periodically and regularly, thus they require more resources to produce "snapshots" of their status quo at a particular moment in time.

Having made clear the major *distinction* between permanent, physical "structures" and ephemeral, dynamic "activities", in the chapters that follow, I try to condense and organize the lessons that have been discussed thus far, to propose a provisional framework for what I call City Knowledge.

In the next chapter, I establish the *premises* for City Knowledge, that revolve around a paradigmatic shift of perspective on municipal information awareness. In essence, I propose the adoption of an "information-aware *modus operandi*" so that towns can begin to treat information as an infrastructural element, as essential as roads, sewers and electricity.

The third chapter introduces the *obstacles* that have hindered the spontaneous emergence of comprehensive municipal information systems around the world. In that chapter, I make a case for why now is the right time to overcome these obstacles and move toward full-fledged City Knowledge.

I then enunciate in the fourth chapter of this Part IV the *qualities* that would be embedded in a comprehensive City Knowledge system in chapter four. Each quality is discussed and examples of how to attain it are provided.

The last chapter sums everything up and distils the six pillars that constitute the *foundations* of City Knowledge, namely the *middle-out* approach, informational jurisdictions, distributed knowledge, sustainable updates, citywide standards, and information sharing.

PREMISES of CITY KNOWLEDGE

Based on my *praxis* as described so far in parts II and III, in this final Part IV, I draw my conclusions and provide a generalized, theoretical synthesis of the lessons we learnt, in order to propose a reasonable pathway that municipalities could embark on to acquire and maintain comprehensive City Knowledge systems.

In the chapters that follow, I introduce the foundations of a methodology aimed at fostering the emergence of a municipal knowledge infrastructure that can be constructed gradually, with a systematic process, without undue stress for local officials and municipal workers, and without breaking the bank. As utopian as this concept may seem, my personal experience has convinced me that this type of "plan-ready" knowledgebase is not only desirable but also quite feasible and sustainable.

single-use, disposable information

Information is already used in cities and towns on a daily basis. Unfortunately though, most of the documentation that is acquired is only used for a single purpose. For instance, the city's building inspector receives an application for a construction permit, with attached drawings and plans that describe an addition to a home. These attachments are only treated as supporting evidence, used exclusively to grant or deny the permit, and then they are shelved and forgotten forever. When the inspector visits the construction site and gives final approval to the finished work, updated drawings are generally not filed away in the city's archives to permanently record the change to the form of the city that just occurred. Unfortunately, in the long run these seemingly insignificant piecemeal changes to individual properties add up to irreversible transformations of our cities and towns. Moreover, failing to record these modifications as they happen has more immediate consequences on the efficiency of the municipal machine. For example, the assessor's department may never receive notice about the increased footprint of the building, therefore real estate taxes for that property will remain unchanged until the next round of appraisals is done to bring the assessed values up to date with market values[345].

failure to record change

loss of tax revenue

When the tax assessors for the town where I live (Spencer, MA) came to visit my house in 2002, the official assessor's map they were carrying was still missing an addition done in the mid 1990's, as well as the more recent addition completed in 2001[346]. I wouldn't personally complain about the fact that they undertaxed my property for more than a decade, but one can see how this lack of attention to information leads to gross inefficiencies that can result – among other things – in loss of revenue.

[345] In actuality, some automated reporting between the building inspector and the assessor does take place, but the inspector does not exchange information with the assessor above and beyond the simple signaling the completion of a renovation or construction. This alert will in turn generate a visit by the assessor. The process is still sub-optimal since the assessor does not receive any concrete information about the work, but merely a notice of completion. In fact, typically no city office receives final drawings or any information worth retaining from the owner or contractor, which represents a missed opportunity for "free" accrual of city knowledge from the ground up.

[346] Together, these additions had quadrupled the living space and more than doubled the footprint.

missed opportunities

Above and beyond these blatant examples, inattention to the value of information is so widespread to be mostly unnoticeable. When a contractor is hired by a city to trim branches off the trees in a certain part of town, nobody is going to notice the missed opportunity to extract an informational return from these services. Indeed, the opportunity for a city worker to be face-to-face with an individual tree is so rare that these periodic trimmings are probably the only chances cities have to get a report on how the tree is doing. As we proposed in Cambridge, MA and Venice, Italy[347], in addition to pruning the branches, these crews could collect numerous pieces of useful, yet simple, information from a brief visual inspection of the trees, and a few quick measurements. For instance, the diameter at breast height and the canopy radius could be measured, so that the city arborist could have an idea of whether the growth of the tree is being stunted or proceeding normally. With minor training, crews could be taught to identify telltale signs of the main diseases, so that a botanist could be sent to the plant for a rapid follow-up to make treatment available as quickly as possible. The workers could also report on the condition of the sidewalks, curbs and storm drains vis-à-vis the trees' root systems, as well as on the distance of the branches from the closest houses, telephone poles and electrical power lines.

informational return

maintenance with a plus

added informational value

These examples of value-added informational extensions to typical city services exemplify the basic tenets of a sustainable municipal knowledge infrastructure. If cities make a conscious decision to extract informational returns from every single activity, City Knowledge will naturally emerge as a byproduct, with only minor additional efforts above and beyond current procedures. As soon as a city adopts an information-aware *modus operandi*, slowly but surely the accumulation of city knowledge will become routine and the shift to plan-ready information will occur almost effortlessly.

information-aware modus operandi

office-wide information assessment

Once the switch to this information-conscious approach is made, in order to be able to adopt this tactic, the various departments of a city ought to reassess their standard operating procedures (SOP) to identify:

- the exact type and form of the information needed to fulfill the various responsibilities of the office;
- the sources of the data to fulfill those informational needs;
- the modifications to the current SOP that could enable the acquisition of information that is long lasting, updatable and reusable[348]
- the extensions that could be made to the current procedures, forms, interactions, and related activities, in order to gather richer information that could lend itself to multiple uses

technical and organizational issues

Of course, there will be technical and organizational issues to be dealt with, but the bottom line is that the transformation of municipal maintenance, management and planning operations will happen if and only if city knowledge principles are embraced at least by one city department as a

[347] See for instance page 128.

[348] We carried out just such an assessment in the spring of 2004, on behalf of the Boston Environment Department. Cf. Hart *et al.*, 2004.

start. Once the collection and organization of information becomes ingrained in any municipal activity, it will be only a matter of time before a thorough and complete knowledge is accrued that will be resilient, reusable, sharable and updateable in perpetuity.

actions to bring about City Knowledge

Once the decision has been made to treat information with the importance it deserves here are two basic activities that need to be undertaken:

1. Collect and organize information about all of the physical elements and human activities that already exist within the municipal boundaries and are either maintained or managed by the city. This is what I will refer to as "the backlog".
2. Develop mechanisms to capture future changes as they happen.

As needed, the update mechanisms could be implemented concurrently with the data gathering, especially since catching up with the backlog may take months or even years to complete. Nevertheless, the backlog will never be taken care of unless some action is taken now. The first and most important move simply consists in consciously deciding to make information as important to the city as other, more traditional infrastructures like roads, sewers and water already are. Such a move would be tantamount to what Thomas Kuhn calls a paradigm shift[349]. When this cultural revolution takes place, the city will be well on its way towards the creation of a sustainable municipal knowledge infrastructure.

information as an infrastructure

a paradigm shift

This new approach to urban ontology seems so natural and obvious that one has to wonder why this paradigm shift hasn't taken place already. The next section explores the financial, technical, logistical and organizational hurdles that have prevented the emergence of city knowledge until now.

[349] Kuhn, 1962.

OBSTACLES TO CITY KNOWLEDGE

The paradigm shift I am invoking is not exactly new nor unique. In fact, what I am proposing may be so commonsensical that it may appear borderline trivial. Certainly, many city pundits before me have proposed bits and pieces of the overall approach I propose – from Geddes[350] to Olmsted[351], from Jacobs[352] to Lynch[353] – and it would be surprising if city workers, who frequently need to make quick decisions in contexts fraught with uncertainty and lack of information, had never wished to have what I call "plan-ready" information at their fingertips. So what has prevented these simple concepts from taking hold sooner? What were the obstacles that have made the cumulative collection of city knowledge infeasible until now?

I think there are at least five main reasons to explain why City Knowledge has not been practical until very recently:

1. the perceived high cost of such an enterprise, both in terms of the initial cataloguing effort and in terms of the subsequent upkeep of the information;
2. the difficulty in cross-referencing different data archives, especially before the advent of the PC and desktop databases;
3. the intricacies of relating information to specific locations in space, even when street addresses are used;
4. the complexity of coordinating and synchronizing data within and across agencies;
5. the frustration that many municipal officers have experienced when trying to keep up with constantly changing technology, even after the introduction of computer tools into municipal operations.

The five "problems" are closely interlinked and the solution of the spatial referencing dilemma thanks to Geographic Information Systems (GIS), together with the ease of cross-referencing brought about by relational database management systems (RDBMS) have combined to greatly lower the cost barrier, making the whole City Knowledge proposition feasible and affordable in the late 1990's.

The biggest problem remains a "people" problem, due to the all-to-human reluctance to accept and adapt to changing situations and technologies[354]. "The people and organizations designing and managing GIS often are uninterested in such comprehensive systems"[355] so, after a discussion of the five problems listed above, I dedicate a final section to this overarching problem of the adequacy of the human skill sets for the tasks

350 Geddes, 1911.
351 Olmsted, 1913. See quote on page 159.
352 Jacobs, 1961. See quote on page 123.
353 Lynch, 1968. See quote on page 39.
354 Reeve and Petch, 1999, p. 5.
355 Innes and Simpson, 1993, p. 232.

required and the consequent resistance to change that such endeavors invariably encounter.

the Cost of City Knowledge

Inventorying the physical urban infrastructure already in existence is an imposing task. Before the advent of personal computers and relational databases, systematic and exhaustive tracking of city assets was very cumbersome and consequently not very flexible. Paper records were maintained (and often still are) in file cabinets, using a variety of ad-hoc indexing schemes[356] suited to the mission of the office where the records resided. Re-indexing and cross-referencing were simply not available options if one wanted to re-utilize an existing archive for practical or analytical reasons that differed from the original intended purpose of the documentation. Enriching the archive with complementary information that augmented the core collection of indispensable data was not even contemplated, given how unwieldy these paper stores were, even when the bare minimum of necessary information was retained.

lower cost of computerization

Nowadays, the widespread adoption of computerized databases for many municipal operations has greatly reduced the cost of keeping the records organized, and of making them accessible for multiple purposes. In fact, some of the bigger towns have even begun to use the web as the vehicle for making the information more accessible to citizens, though a full two-way interaction is still not commonplace even in the more advanced e-government systems in operation[357].

As mentioned earlier[358], there are two principal tasks a municipal office needs to address once the paradigm shift has taken place toward an information-aware modus operandi: (1) create computerized inventories of the pre-existing city structures and activities already "out there" and (2) make sure future change is captured as it happens.

creatively catching up with the backlog

The cost of computerizing the "backlog" of information that is already in our municipal archives – or simply already in existence in the "real world" – may appear to be prohibitive for some communities, but there are numerous ingenious ways to make the process affordable even for cash-strapped municipalities. Some of these creative approaches may include:

1. Making the task of computerizing past records part of the daily routine of some of the municipal staff already on payroll;
2. Leveraging inexpensive (or free) volunteers such as interns, summer workers, university students or high-school students to do the bulk of the data entry;

356 Like "place-over-time" as Bryan Glascock of the Boston Environment Department describes the typical filing system where permits and other paperwork are first of all filed in folders organized by address ("place"). Within each place-indexed folder, one would then find the documents organized chronologically ("over time").

357 Hart *et al.*, 2004.

358 Page 164.

3. Outsourcing the computerization of the pre-existing situation, through contractual obligations, to contractors that are hired for routine maintenance tasks;
4. Actively pursuing "instant gratification" in the form of a rapid financial return by investing on applications that promise to yield immediate economic benefits;
5. Setting aside some funds (5-10%) from the ordinary budget of each department to gradually computerize the backlog of paper records already in the municipal archives. These funds could either pay overtime for regular staff members or could go toward paying some outside consultants to do the job.

information-conscious job descriptions

Perhaps the most cost-effective manner to make City Knowledge principles a reality in any municipal office is to simply decide that information is utterly important to city operations. Once such a momentous decision is made, the city can review the job description of every civil servant in town with the intent of extracting the maximum informational return out of every worker. So, the next time crews from the Department of Public Works (DPW) go out to unclog a drain after a storm, the "new" job description will force them to map (however approximately) the location of the drain with respect to a nearby landmark. The drain will thus not only be unclogged but it will also be recorded and codified so that future interventions will be able to refer to it by its code-name and will also be able to detect patterns of clogging that may have gone unnoticed until information was accrued with such specific place references.

maintenance-based data collection

However slow this process of information recording on a "need-to" basis may appear, it is guaranteed to eventually produce a complete map of all features that are under the jurisdiction of each department. As an added bonus, the data collection will also be done on a priority basis, leaving the less trouble-prone elements for last, as logic would suggest. The drawback of such an approach is that it would require crews to be always on the alert and ready with the appropriate field forms and data-collection equipment. The marginal returns may not be worthwhile once the occurrence of an unrecorded element becomes more the exception than the rule.

free or inexpensive labor

Another low-cost approach to the computerization of the backlog is to enlist the help of *pro bono* volunteers or inexpensive summer interns to do the data collection. In many ways, this is the approach I have personally undertaken through the hundreds of WPI undergraduates and Earthwatch volunteers that I guided through the data collection campaigns I described in previous chapters[359]. Towns already enlist the help of interns for several purposes, so the difference I am proposing would be simply in the sustainability and "staying power" of the work conducted by these volunteers. As with anything else, the ultimate impact of the work carried out free-of-charge by volunteers is only as good as our ability to follow up and use the work – hopefully more than once. Therefore, this money-saving option would entail managing the volunteers in a way that will produce continuous growth of the knowledge database. Making good use of free or

[359] Page 39 and following.

inexpensive labor requires careful management and adequate attention lest it become another exercise in futility. Keeping interns "busy" is not a guarantee that city knowledge will accrue incrementally. Data collection, if carried out by volunteers, is only as good as the methodology employed in the process and is only as usable as one's ability to turn data into information.

contractual obligations

Another surefire way to minimize costs while knowledge is being gradually and systematically augmented is to revise all outsourcing contracts to include informational returns as part of the services rendered. It will be up to the contractors to equip themselves to fulfill the new knowledge-focused contracts, which require updated information in digital form. Thus, the city could demand that the company that is doing the usual pruning of the city trees will also measure the circumference of the tree trunk, estimate the distance from power lines to the nearest tree branch, assess the condition of the sidewalk vis à vis the tree's root system, and conduct a thorough check of visible signs of potential disease. In this vein, the City of Cambridge, based on our suggestion, has obliged its contractor (Lockheed) to keep track of parking ticket information in a manner that will allow analyses to be made to resolve problematic situations[360].

low-hanging fruits

A really convincing way to dispel any doubts about the economic viability of the construction of a City Knowledge system is to look for "low hanging fruit" in the form of City Knowledge projects that will yield instant financial returns, amply justifying any up front outlay of funds. An example of low hanging fruit would be the creation of a catalog of parking meters and curb regulations that would enable parking control officers to be more efficient in their routes[361]. Improvements in efficiency would greatly increase the ability in detecting infractions, which would in turn yield immediate and permanent returns quantifiable in tens of thousands of dollars.

budgeting for city knowledge

My experience suggests that espousing City Knowledge principles saves money – if nothing else for the demonstrated reusability of the datasets. Once a sufficient number of City Knowledge projects will be operational, a more exhaustive economic analysis of the advantages of these systems should provide conclusive evidence to prove or confute this inductive assumption. Eliminating redundancy should at the very least free up lots of time to focus on more fun endeavors than the hunting down of datasets to analyze. A careful weighing of the pros and cons may move a city to simply decide to invest in the creation of an emergent, self-sustaining, once-and-for-all city knowledge system. A yearly amount budgeted for such a cause would initiate the emergence of City Knowledge and would be well worth the investment even if the value-added synergies that I described in previous chapters should fail to materialize. Having well-organized information at one's fingertips will make all operational decisions more

[360] Flynn *et al.*, 2003.

[361] As in Flynn (2003) when our team was able to suggest practical procedural changes to the City of Cambridge, that enabled parking control officers to warn parking meter collection crews about jammed meters on a daily basis. Hundreds of thousands of dollars could be saved with this simple procedure.

effective and efficient[362]. Eliminating redundancy is the minimal result one is going to obtain from embarking in a City Knowledge path and such an outcome is incontrovertible and guaranteed.

the real cost of city knowledge

Inventorying all 472 Venetian bridges cost Insula about 74,000 dollars (€61,000), i.e. $ 156 per bridge. Similarly, our survey of 1627 docks in the entire lagoon cost the city $250,000 (€205,300), which translates to $153 per dock, which is amazingly identical to the per-bridge cost[363]. The savings derived from these multimedia catalogs are hard to quantify, but the cost per object translates into roughly 10 person-hours of work[364]. Would more than 10 hours be spent by someone gathering these same data from now until the end of time? If so, how soon would someone have to come along to collect such data? How often would such a redundant activity take place? These questions are hard to answer definitively. We spend much more time and money on finding, requesting and obtaining data than on manipulating the data and analyzing results[365]. So, I personally think that a city or town that was prepared to invest in City Knowledge would be making a smart choice. If the will to purposely budget for City Knowledge is not there, then one could still resort to the other low-cost or no-cost solutions described earlier.

the cost of knowledge maintenance

While the backlog is being whittled away, one needs to worry about the maintenance of the knowledge being acquired. Knowledge updates always represent a cost for someone – either in terms of money or in terms of time. From the perspective of a municipality, the trick is to extract as much information as possible from the private sector without significant cash disbursements. As long as we restrict the discussion to the maintenance of information about "things", it is possible to imagine how maintenance of the physical object could be coupled – by design – to activities aimed at revising and/or verifying the underlying dataset. As mentioned, the company that is contracted to prune all of the city trees would also be charged with measuring the circumference of the trunk, the height of the tree, the shortest (and more "dangerous") distance between branches and power lines and the canopy radius. These scheduled maintenance activities would thus contribute to a periodic updating of the underlying knowledge-base as well. Eventually, the wealth of up-to-date knowledge available to the city arborist would perhaps dictate a different scheduling or sequencing of the physical maintenance operations, allowing for an optimization that is far from possible today.

self-reported updates

One "free" updating strategy that worked for the Boston Air Pollution Control Commission (APCC) was to let the "customer" (parking facility owner), do the updating of much of the essentials, by instituting a periodic renewal of the permit or license that required the submission of

362 Budić, 1994, p. 252.

363 This amazing agreement between such disparate projects is worthy of further investigation, but if provides a wonderful heuristic for ball-park estimates.

364 The about $15.00 per hour used here is an approximate average of the hourly cost for the more skilled tasks such as GIS and database manipulation, together with the less demanding tasks of field measurement and surveying.

365 Budić, 1994; Nedović-Budić, 2000, p. 82.

updated statistics[366]. In general, self-reported data need to be validated, but linking the submission of information to fees may, in some cases, provide a highly reliable means for information maintenance, while also providing funds for the upkeep of the datasets[367].

administrative updates

In general, a focus on informational returns has yielded many surprisingly simple ways to acquire updates with minimum effort, as was amply discussed earlier in this document. By far, the most promising way to keep the knowledgebase up-to-par – insofar as change is produced by human acts and not by natural dynamics – is to actively work to intercept and process the administrative paperwork that accompanies such change, since almost all anthropogenic modifications to the world we live in are decided, requested, required, approved or authorized by some level of government.

Disconnected Datasets

Until the advent of the personal computer in the 1980's, it was inconceivable for city officials to even imagine how different archives could be cross-referenced to produce augmented information from mere documentary data. Only a handful of hard working scholars would ever attempt to do something like that even in a very limited research domain since such a project entailed literally consulting hundreds of paper files, trying to reconcile them with each other to glean at some hidden pattern that explained some interesting phenomenon. It was hard enough to thumb though a single paper archive, never mind two or more. Ironically, it was easier to do such arduous research on very ancient materials than on very recent ones. The paucity of antique records made the task more manageable than it could ever be in the presence of miles and miles of massive shelves brimming with modern paper records.

relational databases vs. flat-files

The arrival of the computer age in the 1970's did little to improve this situation, since the early tools were first and foremost geared towards the "keeping" of the records and not so much for their analysis. Eventually, in the following decade, PCs began to make their appearance on the desktop of researchers and scholars and the first personal database tools became available to the masses under the novel operating system nicknamed DOS created by an upstart computer company called Microsoft. Ashton-Tate's Dbase III was the first database that gathered a certain following in the DOS community. The "relational" capabilities of such database tools were hardly ever tapped, however, and even today "flat-files" seem to prevail, which prevent datasets from being connected into wholes that are bigger (or at least more informative) than the sum of their parts.

disconnected vs. distributed

Paradoxically, this "insular" mentality was further aided by the subsequent mass diffusion of the personal computer in the 1990's, which was not accompanied by a parallel dissemination of networking hardware and software. Stand-alone applications held sway until the end of the

366 As we suggested in Allard *et al.,* 2001.

367 In the case of the APCC, a yearly fee per parking spot was instituted so the owner would immediately report spots that were no longer available for public parking. Of course, this system could be vulnerable to gross underreporting, but at least we eliminated the opposite problem of over-estimating parking availability. See also Eichelberger, 2004.

nineties when finally the usefulness of networking became more apparent to everyone and connecting PCs became cheaper and easier, as the World Wide Web became an overnight sensation since its modest beginnings in the early-nineties[368].

Creating an emergent system that would grow knowledge from many different-yet-connected systems, managed and maintained by different agencies having jurisdiction over different urban domains, only became truly possible very recently. Until the end of the last century (and millennium) such a system was a mere utopia, though the technical difficulties involved in such an endeavor did not stop some visionaries from predicting the day when such a networked, distributed intelligence would be commonplace[369]. I was fortunate to be there during such a momentous period of human history and I have kept in my personal archives documents that I wrote in the 1980's that hint at such a system. In my own way, I was one of those visionaries…

organizational disconnects

Today, technology has finally made good on the promises of those heady days[370] and the difficulty in cross-referencing different data archives is organizational and no longer technical. It is now people, departments, organizations and agencies that create the barriers between datasets that make a truly emergent city knowledge system difficult to implement[371]. This dissertation is my contribution towards the breaking of these artificial barriers that prevent such a distributed knowledgebase from being widely used to maintain our urban infrastructures, manage our civic activities and plan the future of our cities.

Vague Spatial References

Another major obstacle to the widespread adoption of the intuitive mechanisms that constitute the City Knowledge approach has been the inadequate manner in which our municipal recordkeeping has dealt with references to geographic locations. Once cities began to keep track of data using computers, their primary purpose was to manage financial records, such as payroll, taxes, fees and fines. To this date, spatial references – when present at all – are limited to traditional street addresses, which, while a step forward, have amply demonstrated their inadequacy for analytical and management purposes.

Even in today's most advanced GIS efforts, buildings rarely referenced by unique IDs[372] and, except for Venice, I have yet to see a city which has coded each individual doorway with a unique code to replace (or

368 For a good on-line recounting of the evolution of the web and other related technologies, see for example, http://www.netvalley.com/intval_intr.html (accessed 6/27/04).

369 See for instance Budić, 1994.

370 McFall, *ENR*, New York: February 16, 2004. Also, Budić, 1994: abstract p. 244.

371 Nedović-Budić, 2000, p. 82; Nedović-Budić and Pinto, 1999, p. 60.

372 The British TOIDs are a step in the right direction, though I find them a bit too arbitrary and not sufficiently mnemonic to be used successfully by humans when needed, altough Building IDs may be one of the few codes that might as well be numeric sequences since it would be hard to come up with a mnemonic identifier.

at least augment) the inadequate yet typical "street and number" addressing scheme[373].

GIS and spatial references

This dearth of geographic references in city-owned datasets has made it difficult to integrate the information available in different departments and has therefore impeded the development of a distributed City Knowledge system. The advent of GIS in the eighties began to change things a little bit, though real progress has been slow since the power for spatial analysis that GIS provide has been misunderstood and underused in favor of more mundane uses of these powerful tools as glorified mapping and plotting applications serving the needs of planning commissions and the like[374]. This underutilization of GIS is a pity, but it has had the beneficial effect of at least making the tool a household word in most mid-to-large-size communities in the developed world[375].

space as the glue for the urban puzzle

Just as networking has only recently come of age, GIS has also reached a critical mass in terms of its widespread adoption in municipalities worldwide, making this decade ripe for the next step, namely the final tapping of GIS's real power as a tool that will enable disparate pieces of the municipal information puzzle to be glued together through the space that they share. A better appreciation of the "power of space" and a better and more educated application of City Knowledge principles for the unique coding of objects in the real world and the linking of maps to databases will enable the emergent qualities of what I am proposing to be unleashed so that the benefits of city knowledge can become evident to all.

Asynchronous Co-dependence

Once we put in place a distributed system that exploits spatial relations and employs well-designed codes to connect pieces of information under different departmental jurisdictions, to create a whole that is superior to a mere compendium of multiple databases, we run the risk of creating dependencies between datasets (and hence between departments[376]) that may spell the ruin of our distributed system. If one piece of the puzzle were to fail, it may take the rest of the system down with it. The fear of such dependency has prevented even the best-intentioned municipalities from embarking in the creation of networked systems, linking the different departments along functional lines. The biggest of these dependencies occurs when the entire municipality is expected to connect into a top-down mega-system that encompasses all of the different departments that feed into a huge central repository. If the big, all-powerful central system fails, nothing works. Fortunately, these monstrous systems are a thing of the past since they have demonstrated weaknesses that have prevented them from ever being implemented in full, thus avoiding the risk of catastrophic failure by simply failing to come on line in the first place[377].

top-down dependencies

373 Though I am sure that there are quite a few other cities that must have done the same. The overwhelming majority probably has not.
374 Budić, 1994.
375 ICMA survey, 2002.
376 See for example Nedović-Budić and Pinto, 1999, p. 56.
377 Keating *et al.,* 2003. See also Reeve and Petch, 1999, p. 5.

bottom-up independencies

The alternative until recently has been a bottom-up trend toward stand-alone computing. The various municipal departments that relied on a number of disconnected PCs and applications to support their work quickly discovered the need to connect, and the corresponding difficulties in coordinating and synchronizing the efforts of different offices[378]. Such apparent complexity discouraged most cities from even attempting to coordinate the various divisions and even internal interactions within a single department were not so common. Such had been the history of "distributed" computing in city governments until the 1990's and the coming of age of the internet.

the interconnectedness of a webbed world

middle-out to the rescue

With the advent of the web, a culture of interconnectedness and a certain familiarity with the concept of sharing through a distributed network of independent computers have created the right mindset upon which the City Knowledge concept of "middle-out" can now be grafted. Middle-out entails that each department will first and foremost take care of its needs, so that the primary functions that the department or office performs will be invariably performed with or without the connection to the outside world. With proper safeguards, each branch office would be capable of functioning on its own, regardless of the state of other offices in the city. Nevertheless, if one department requires knowledge of some aspect of the city that falls under another department's jurisdiction, a City Knowledge system would expect that such knowledge would be shared and that the information would be kept up to date by the department in charge. In a worst case scenario, old-fashioned means of communicating information between departments could be employed and the last-best-version of a dataset could be used if the absolute latest is somehow unavailable at the time.

CK on the information superhighway

Barring the occasional server that goes down, the City Knowledge system discussed herein would rely on normal internet technology that has reached a high level of reliability and resiliency, so the distributed, interconnected City Knowledge infrastructure would be no more vulnerable to co-dependency than our email system is, upon which we already rely rather heavily to conduct our daily business. If something happened that disrupted these systems in a major and long-lasting way, we would be probably facing problems that are much bigger than the mere malfunctioning of our City Knowledge system.

If there are no hardware problems, the real hurdle will remain the difficulty in coordinating the efforts of different departments who need to share some of their information. This problem will not go away magically thanks to City Knowledge, but I think the gradual self-generated transitions that are envisioned in City Knowledge will make these interconnections more likely to be successful than with any imposed-from-above solution.

[378] *Idem.*

Technological Obsolescence

A final concern that has traditionally been cause of a certain reluctance to adopt new technologies in general, and has thus prevented City Knowledge from emerging as a municipal *modus operandi* until now, is the fear of technological obsolescence. As computers became commonplace in the last two decades of the XX century, people have witnessed also a relentless escalation and evolution of both hardware and software. Every year a new improved model or version is released in a never ending race to the ultimate computing power according to Moore's law[379]. In this wild race, many people, companies and municipalities have got burnt at least once by a technology that did not deliver what it promised and resulted in an irretrievable waste of time and money. These bad experiences were so common that terms such as "vaporware" were coined to capture some of the disappointment people experienced when products, companies and entire technologies came and went at lightning speed, leaving behind lots of disgruntled victims, who were thus turned from enthusiastic adopters to cynical neo-luddites. Even the more savvy and fortunate had to endure repeated migrations of their data from one platform to another, having to reinvest time and money at regular intervals to avoid losing years of work by being painted into a technological corner from which there was no escape.

getting burnt by vaporware

growing pains

When we started the project center in Venice in 1988, we used 8086 PC's with 5¼" floppies. Our database was Dbase III; we used Lotus 123 for graphs and – being way at the forefront of technology – we even had Mapinfo for DOS version 1, a GIS program that had just been developed a year or so before by a small company that I visited when both of its workers shared a small cubicle in Troy, New York. We were hugely proud of what we were able to accomplish with those tools. It was already a big step up from the Commodore 64 on which I wrote my undergraduate thesis. Since then, our database migrated to Dbase III+, then IV, then FoxPro, then the early versions of Access up until the current version, and even up to SQLServer and Oracle. Some of the data we use today, though, is still the same we collected way back when. It was not easy, and it was frequently unpleasant and frustrating, but we lived through these transitions with only a few scars to show. We wasted thousands of dollars in bogus hardware and software or on products that served us for only one season before being discontinued. But I would do it all over again because, despite all that, we inched forward and finally blossomed in the mid-nineties, tempered by our harrowing experiences and all the better for them.

future evolution through migrations

I would be delusional if I were to suggest that the internet and GIS and databases as we know them today are the ultimate tools that a city will need to use in order to maintain its data for posterity. Many amazing technologies will come and go in the years to come and what today seems utterly unbelievable and totally awesome, will some day (not too long from now) seem silly, childish and banal. Nevertheless, the mechanisms, the

[379] In 1966, Gordon Moore suggested that the number of transistors in a microchip would double every 18 months. This original quote was subsequently extended to cover the doubling of computing power and halving of price every 18 moths (see http://firstmonday.org/issues/issue7_11/tuomi/).

structures, the procedures, processes and codes, the jurisdictional partitions and the overall City Knowledge approach that I propose should have staying power way beyond any technology that we use today.

What counts here is not the tool but the approach. Data can always be migrated to a new format as the last 15+ years of my personal experience can prove. However, badly-structured data will remain bad, whereas a well structured dataset and an accurate map layer will remain. If the data and layers relate to permanent (or very slowly changing) features of our urban landscape, then we can be sure that whatever effort we put in today will not go to waste because of the vagaries of technological advancement.

Resistance to Change

The biggest obstacle for the development of tools to deal with significant planning tasks that "require comprehensive, multipurpose, and multiuser geographic information systems" have been people[380]. Resistance to change is in some measure due to self-perceived inadequate skills and lack of training programs to ameliorate them. Another part of the resistance that planners experience is probably imputable to the separation between planners and the repositories of data[381]. The "fear of losing autonomy, control over information sources, independence, and organizational power is widely acknowledged"[382]. There seems to be a consensus among researchers that planners have a "limited vision of the potential of GIS[383]" and this has resulted essentially in a stagnation in the development of Planning Support Systems despite the great technological advances of the last decade[384].

Although, "the most important impediment to the implementation of GIS in planning may be planners themselves[385]", in this paper I propose a distributed system of data accrual and sharing that may allow planners to skip completely the issues of data collection that take up so much of their time[386] and thus be able to focus on the more challenging issues that planners are supposed to concentrate on. The key to overcoming the "issues of organizational inertia, mistrust, and "turf[387]", according to Innes and Simpson[388] is to follow an implementation path that displays the following traits: simplicity, observable benefits, relative advantage, ability to make small trials, and compatibility. I think that my City Knowledge approach incorporates all five of these principles.

The preeminence of this obstacle is why I think that the most important step toward City Knowledge is to accept information as a core component of the city' infrastructure, on par with water, sewer, roads and electricity, and to begin treating it as such in all aspects of municipal

380 Innes and Simpson, 1993, p. 232; see also Reeve and Petch, 1999, p. *xiv* for ex.
381 Klosterman, 2001, p. 4; Reeve and Petch, 1999, p. *xiii* and p. 6.
382 Nedović-Budić and Pinto, 1999, p. 54.
383 Innes and Simpson, *idem*.
384 Klosterman, 2001; Geertman and Stillwell, 2003.
385 Innes and Simpson, *idem*.
386 Nedović-Budić, 200, p. 82.
387 Nedović-Budić and Pinto, 1999, p. 60.
388 Innes and Simpson, *idem*.

operations. This paradigmatic shift alone will generate the rest of the transformations needed to gradually bring City Knowledge to be embedded in the municipal *modus-operandi*.

QUALITIES OF CITY KNOWLEDGE

In this chapter, I present a sort of "wish list" that contains all of the positive qualities that a City Knowledge system should display. After explaining why each quality is desirable, I point to specific instances from my personal experience or from the literature that demonstrate that such a quality is indeed achievable in a municipal information system and I then discuss how that quality can be achieved by a City Knowledge system. I this chapter, I will try to address each of the obstacles from the previous chapter, while also referring back to the lessons sprinkled around Parts II and III to support my arguments.

In my opinion, the distinguishing qualities that a comprehensive City Knowledge system should try to achieve are[389]:

- Affordable and Easy-to-assemble
- Gradual and Systematic
- Permanent and Exhaustive
- Sustainable and Up-to-date
- Rich and Reliable
- Flexible and Re-usable
- Shareable and Secure

AFFORDABLE

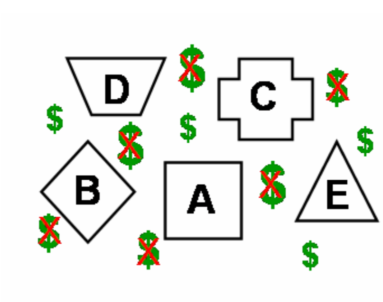

lower transaction costs

Capturing the information for all of the structures and activities that are already "out there" in an urban environment may seem like a daunting task. Catching up with the backlog will have a cost associated with it, but the expense can be amortized over a long period. As shown in previous chapters, most of the data we collected in Venice was gathered by students working *pro bono* (actually *pro grade*). In a typical year, we only had 24 students in Venice, for a period of only two months, yet we were able to acquire an impressive collection of datasets on a variety of different aspects of the city. Though there are costs associated with these endeavors, they can be defrayed in a variety of creative ways, as explained in the previous chapter.

The archival of city knowledge that I propose makes practical and economic sense today due to the declining costs of such activities, thanks to the technological evolution that has brought databases and geographic information systems (GIS) into the mainstream of municipal operations, even in smaller towns. Once the existing state of things is recorded and organized in databases and GIS layers, the task will then be to intercept change on a day-to-day basis, so that there will never be the need to catch up with backlogs again in the future. If done carefully, this constant upkeep of information should cost very little additional money. The difference between the current procedures and those that will be put in place in the

[389] I think that all of the five principles for an effective implementation of GIS that Innes and Simpson (1993) list (see page 175) can be mapped onto the qualities discussed here.

future is really minimal in terms of resources. The main difference is one of focus, as discussed in the previous chapter[390].

Easy-to-assemble

As already mentioned, we are at a peculiar juncture in the advancement of information and communication technologies (ICT) that enables us to take for granted computer tools that simply were not widely available as recently as the late nineties. Each individual element of city knowledge is typically not that complicated to capture in databases and/or geographic information systems. In fact, one of the contentions that is made in a later section is that the fundamental elements of the urban realm – from trees to streetlights, from roads to traffic lights, from park benches to public art – are rather common across the international municipal landscape.

uniform characteristics

The universal nature of the components that make up a city creates the possibility of the development of standards that parameterize the uniform characteristics of typical city assets. Such a solid foundation will then enable supplementary customizations to suit the peculiar needs each specific township. The relative simplicity of the parameters that characterize each category of objects (or of actions, as in the case of traffic) makes it possible to train staffers, volunteers and contractors in the procedures necessary to consistently collect the information in the field[391]. A fundamental tenet of our approach is to always atomize the parameters so that each aspect is gathered in a manner that is as objective as possible. Our data collection always relies on visual inspection and simple measurements, with the occasional support of more sophisticated instruments when necessary. Once again, the proof of the fact that these activities are easy to conduct lies in the fact that our massive city knowledge effort in Venice was carried out by twenty-year-old students, who did not even speak the local language, yet were capable of gathering all of the necessary data even about items that they had never heard of before their projects began.

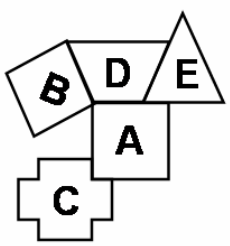

Gradual

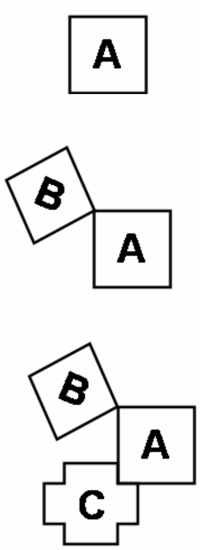

It took us more than a decade to complete some of our largest databases in Venice. As mentioned, tens of thousands of student-hours went into the creation of our public art and canals information systems. Yet, in any given year only a handful of students devoted a maximum of just two months to each of these undertakings. The trick for us was to tackle one borough at a time, so that at any one point we would always have some parts of the city completely done. The difficulty, from the academic and pedagogical perspective, was to propose ever-challenging projects, even though the topics may have been the same as the prior year's, albeit in a different part of town.

The way around this conundrum was, on the one hand, to focus on the analytical aspects of the projects, challenging the teams to higher-order and more complex analyses of the data, thus making the data collection just an incidental part of a more sophisticated study that tested the critical-thinking abilities of our students. On the other hand, once all of the challenging analytical angles were exhausted, as a responsible educator I had

---

390 See page 166 ff.

391 See for example our tree projects in Cambridge (Creps *et al.*, 2001) and Venice (Bennett *et al.*, 2001).

to stop proposing projects for my WPI students on that specific topic. I made up for this by continuing the more mundane data collection work using volunteers. Much of the public art catalog was completed by dozens of pro bono assistants under my guidance as Principal Investigator of research projects funded by an organization called *Earthwatch*[392]. Volunteers, interns, gradeschoolers and university students all can play a role in the gradual accrual of city knowledge, as they did for us in Venice.

[better late than never]

The key here is to accept the fact that since these objects have been "out there" for decades (centuries in the case of Venice) without being thoroughly investigated and inventoried, it won't be a problem if we take our time cataloguing them now. Even if it takes a few years, they're not really going anywhere and a slow progress is better than none. Conversely though, if we do not begin the process now, we are guaranteed to never see it done.

All of the structures and activities that make up our urban reality seem so intricate, complicated and innumerable that it is hard to fathom how we could hope to actually keep track of all this complexity. The principle of *graduality* is one way to deal with the apparent immensity of the task. Biting just what one can chew is a good way to "divide and conquer" the apparently insurmountable hurdles.

392 www.earthwatch.org

SYSTEMATIC

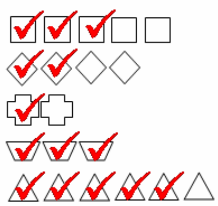

To guarantee that the final product is complete, though, one needs to be *systematic* about the way in which data are collected. In dealing with the backlog, we need to be deliberate and methodical, because the records of past changes will become permanent features of our municipal information infrastructure. Since we waited so long to get this "city stuff" organized, we might as well take our time gathering the information so that we don't miss anything along the way. In Venice, some of our projects benefited from the ancient city's idiosyncratic addressing system, which allowed our students to systematically visit every single address in each borough, in sequence, to make sure that nothing was missed along the way[393]. Similarly, in Boston and Cambridge, target streets were combed thoroughly by our teams to ensure full coverage of the particular aspect of city knowledge being studied.

The modular partitioning of the city into manageable neighborhood units that were analyzed one at a time is another way to ensure steady progress toward the completion of the full inventory of pre-existing urban elements in a *systematic* manner.

PERMANENT

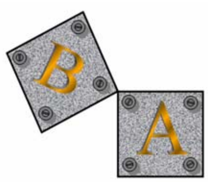

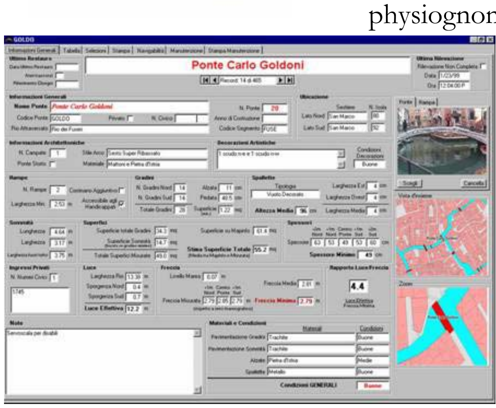

permanent vs. dynamic datasets

Cities may be seemingly intricate and unwieldy, but their physical make up changes really slowly. If we go back to Kevin Lynch's simple separation between structures and activities – the container and the contained – we can see that much of what makes us think of cities as very dynamic and ever-changing is due to the frenetic pace of activities that take place in their streets and sidewalks. Despite all of the observable vitality, the structures that provide the backdrop for those activities remain practically unaltered from day to day. The elements that make up the concrete physiognomy of our hometowns are finite, enumerable and thus eminently recordable once and for all (or *una tantum* as we would say in Italy). It is my contention – borne out of my experience – that it is possible to capture all of the pre-existing features of our material urban environments in a gradual and systematic way and thus to create permanent records that will make it unnecessary to go back to collect any more information about these objects ever again[394]. Realizing the immanence of the tangible city makes any effort at permanently recording city knowledge worthwhile. While it is hard to quantify the exact financial benefits of the efforts I propose, it can be logically argued that a one-time-only, in-depth, systematic and gradual campaign to organize the information about municipal assets for perpetuity will save money in the long run, vis-à-vis the costly and often redundant consultant studies that are commissioned year after year to collect data demanded by the "plan *du jour*".

It must be remembered, however, that even the most immutable features of a city will always be subject to changing conditions, which is why all of our databases are always separated into permanent and dynamic

393 In the summer of 2004, for instance, a team of students followed the entire address space systematically to map all of the storefronts in Venice.

394 This statement is not entirely correct since even "permanent" data will change ever so slowly, but for the purpose of this discussion, such splitting of hair has been purposely omitted.

components for each asset. For instance, bridges in Venice have been there for almost a thousand years, so one database was dedicated to their invariable characteristics, like the span, number of steps, height of rail, clearance, materials and other such features. According to the classic "entity-relation" model for Relational Data Base Management System (RDBMS), alongside this permanent dataset we created a dynamic database – linked to the former via the unique bridge code – that captured the physical conditions of each bridge at the time of our inventory, listing such things as the damage to the steps, the physical integrity of the arch, the state of conservation of the pavement and many other aspects affected by wear and tear over time. Insula, the company that commissioned our bridge inventory, later proceeded to create additional linked databases to keep track of maintenance work conducted on each bridge, which in turn lead to the updating of the time-stamped condition assessment database records.

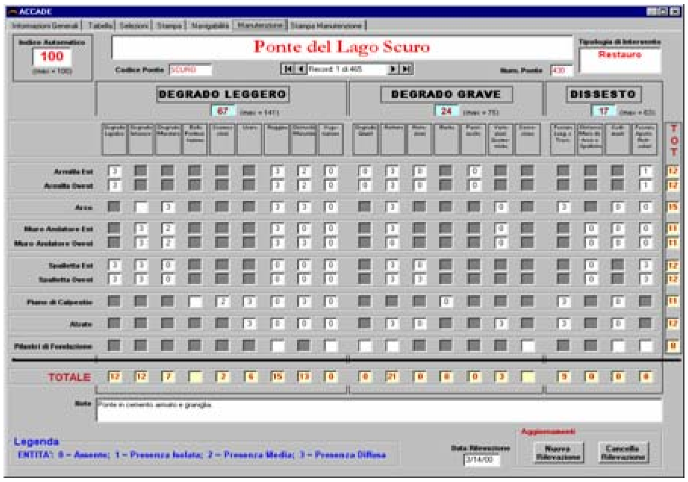

EXHAUSTIVE

informational equity

A gradual and systematic approach to the collection and organization of permanent urban features will eventually result in a comprehensive municipal knowledge infrastructure. Unlike the happenstance consultant reports that focus on this or that part of town as the need arises, the approach I propose will not leave any neighborhood out of the picture. The availability of information about specific areas of a city will not depend on the vagaries of past studies, but will be guaranteed to be *exhaustive* and complete for every single borough, block and street, with no exceptions[395].

In some ways, this feature of my proposed municipal knowledge infrastructure will address issues of social equity that are often hidden in the confusing piecemeal approach to urban information that currently prevails in cities all over the world. It may be argued that the current state of affairs is so fragmented and disorganized that there is an odd form of equity at play in that every part of town and every social class is equally subjected to the inefficiencies that are ubiquitous in our municipal governments. City knowledge, as I see it, will correct this oddity and ensure that the right kind of informational equity is attained through plan-ready information.

SUSTAINABLE

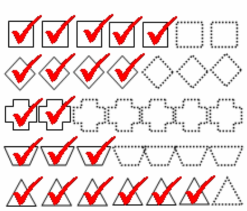

Cities may be slow to change, but they do change a little bit every day. Indeed, it may be a slight misnomer to call some of our urban information "permanent" when in fact it is just "changing very slowly". At any rate, complete city knowledge is a moving target that requires constant upkeep. For a municipal knowledge infrastructure to be truly sustainable, it needs to be constructed in such a way as to facilitate updates in the most cost-effective, transparent and effortless manner. To avoid cost overruns,

395 Our completed Public Art catalog in Venice (see picture on page 110), with over 4,000 records, for example, is exhaustive of all public art of each type.

the data that are collected to feed this urban knowledgebase must be consistent with what is typically acquired by cities in their normal operations[396].

Maintaining the data should be a "natural" extension of typical current practices. The updating effort should be commensurate with the usefulness of the information. To be true to the tenets of sustainability, we should not go out of our way to willy-nilly gather data that expends resources that we are borrowing from future generations of citizens, particularly when the effort is costly and the returns are only marginal. A sustainable municipal knowledge infrastructure does not live beyond its means nor does it waste precious informational capital.

informational capital

The resolution of the data included in a City Knowledge system today needs to only address today's needs and can only be commensurate to the current technological capabilities. A sustainable City Knowledge system will adapt to changing circumstances[397] and to evolving technologies, so it needs to be resilient. All of the qualities listed herein contribute together to making these systems adaptable to varying degrees of precision and data quality.

Up-to-date

updating information about "structures"

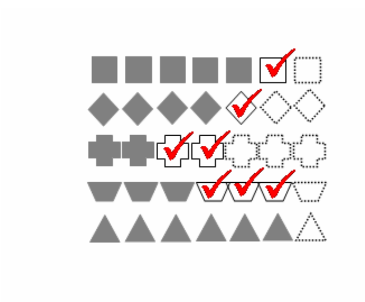

[transactions]

Updating the portion of city knowledge connected with the physical structures that change only sporadically should be simply a matter of intercepting administrative transactions as they happen within the already established procedures that are part of standard municipal operations. This means that from now on, when an inspector gives the "thumbs up" to a newly constructed addition, data from the corresponding drawings are entered into a "buildings database" and, on the GIS side, the "buildings layer" is updated to reflect any change of footprint. The more sophisticated municipal systems may even keep track of internal volumetric changes, by linking 3-D CAD drawings to the building records[398].

These procedures could be institutionalized in such a way as to cost next to nothing to the local community, by simply making the owner/contractors submit these attachments according to predetermined formats[399]. Notice of these changes could be cascaded down to the assessor's office immediately to spark a new appraisal that would in turn be parlayed into an updated tax bill. As unpleasant as such efficiency may sound (especially for those of us who own a home and pay real estate taxes), these revenue-generating examples are meant to emphasize the measurable instant returns a town could obtain on rather modest investments aimed at streamlining current operations.

updating information about "conditions"

Even though the fundamental character of urban structures does not change dramatically over time, their *condition* or *state of conservation* changes constantly since most of these items are outdoors, exposed to the elements.

396 See Innes and Simpson, 1993.

397 As we had to do when we constructed the boat traffic model in Venice, which forced us to add pseudo-nodes to the canal network wherever there was a dock or a bridge.

398 I have advised a Masters thesis in Civil Engineering at WPI on precisely this topic (Samdadia, 2004). Michael Batty in the UK and many others have also been pursuing the integration of 3D CAD with GIS.

399 As we proposed in Cambridge (Gage, 2003).

For some urban elements – like city blocks or parking garages – talking about state of conservation does not make much sense. For some other city assets, condition assessments may need to be conducted only occasionally – as would be the case for traffic signs, for instance. Some categories of physical objects, however, may require frequent updates as do trees for instance, which may become infested by parasites if neglected for too long, or may endanger people or property with their limbs, especially in bad weather. Anything that may constitute a hazard for public safety usually gets priority status and is monitored more closely. Worsening conditions usually entail physical decay and potentially dangerous static deterioration of the object, with perilous consequences for people and properties near the object.

[transactional snapshots]

[citizen complaints]

Fortunately, the worse states of decay are usually reported to city hall by concerned citizens who thus perform this surveillance duty for free in order to protect themselves, their properties or their loved ones. Properly managing citizen complaints can be an effective way to inexpensively monitor the most delicate and treacherous deteriorating conditions around the city. Provided these reports are followed by some corrective action, some form of citizen *vigilantism* may be cost effective for the municipal treasury as well as empowering for neighborhood communities. It is worth remembering that a well maintained city is respected by its citizens, as mayor Giuliani's policies proved in New York city in the 1990's. Enforcing maintenance standards in public and private property can create virtuous cycles of overall improvement of the quality of life in a community, dramatically contributing to the reduction of petty crimes against property and even diminishing the frequency of more serious felonies in the long run[400].

updating information about "activities"

Keeping current with the information about the "activities" within a city's boundaries requires an approach that is different from that used for tracking permanent physical characteristics and is also somewhat different from the methods one can adopt to monitor changes in the state of conservation of these material elements of the urban environment. Dynamic processes, like traffic flows, or business vitality, or demographic change are harder to maintain up-to-date by simply tracking the administrative "red tape".

[snapshots]

[administrative proxies]

Keeping track of new car sales or new car registrations may provide proxies for some of the dynamics that urban managers and planners are interested in, but they will never tell us where a car customarily travels to, or at what time or with what frequency. Toll-road accounting systems and other automated devices can fill the gap, though privacy safeguards prevent us from using the data at the finest grain that is technically available, as we experienced with the Venice boat-monitoring projects described earlier[401]. Although there is some unexplored potential for a more careful analysis of how existing record-keeping systems may help us in developing a framework for the monitoring of these activities, the bottom line is that these dynamic

[automatic devices]

400 Gladwell, 2000, pp. 144-146.

401 Starting on page 86.

practices need periodic monitoring to provide the needed information to administrators and decision-makers.

Automatic devices, such as the ones described in earlier sections[402], can be brought into the system to collect reliable data 24/7, greatly reducing the costs of data collection, after an initial investment in the required hardware and software.

[grain and privacy]

Different agencies customarily collect information about ephemeral activities with reasonable frequency[403], though these informational snapshots often present shortcomings of one type or another. Opportunities exist to leverage existing procedures to rein in some of the possible "free" updates that may be available to municipalities, as was discussed in our experience with traffic data in Cambridge, Massachusetts[404]. Another creative and educational way to conduct updates is to involve local schools. For instance, in Venice we enlisted the help of local gradeschoolers to "keep an eye" on the public art collection by promoting yearly "treasure hunts" for the school children, to ensure that the artwork was still there and in good condition year after year. Beyond that, city planners should make sure that adequate funds are available to collect whatever datasets are deemed essential about these more short-lived processes that take place within our cities and towns.

Rich

While we need to be sustainable in our practices and limit our data collection to the realm of established procedures and to facets of urban life that merit attention and are already acknowledged as important to city maintenance, management and planning, we also need to make sure that we do not flatten our data accrual in such a way as to make it impossible to turn our datasets into a re-usable and sharable information infrastructure.

horizontal sharing

vertical aggregation

Our base data need to be *rich* enough to allow both horizontal sharing and vertical aggregation. If the fundamental datasets are too plain and/or too specific to a particular task, they will not lend themselves to multiple uses, thus they will not allow municipalities to enjoy the economies of scale and savings that could be obtained when the same data are reutilized in a different context without the need for additional data collection[405].

fortuitous vs. intuitive vs. teleologic

The richness of a dataset cannot rely on fortuitous coincidences that somehow make it possible to take data from one particular realm and use it in an unforeseen way to analyze another facet of urban operations. These happenstances may give us useful insights, but we need more than that to create a sustainable knowledge infrastructure. In my work in Venice, I relied on intuition to enrich the data that my students were collecting, by suggesting the inclusion of parameters and measurements that were not immediately useful to the task at hand, but were collected "just in case" to make possible some ulterior use of the dataset at a future date in some

402 See page 83 ff.

403 See, for example, the efforts by the National Neighborhood Indicator Partnership (NNIP, http://www.urban.org/nnip) and by Neighborhood Knowledge Los Angeles (NKLA, http://www.nkla.org), to name a couple.

404 See footnote 317.

405 Specific cost-savings are reported in the literature by Nedović-Budić and Pinto, 1999, and others. The re-usability of the data can be gleaned from the example sin Parts II and III, like for example on page 66.

foreseeable situation down the line. This intuitive approach needs to be translated into a more cognizant teleological method of data enrichment that makes the foreseeable future applications much more overt and hence makes the dataset extensions more explicit and codified[406].

limited "face time"

One needs to remember that "face time" with the physical objects (or activities) out in the urban domain is a rare commodity and these infrequent opportunities need to be exploited for the maximum benefit[407]. When a city worker is sent out to change a light bulb on a streetlight, this ought to be seen as an occasion to get an update on the condition of the entire light post (Does it need painting? Is it damaged? Are the cables hanging low? Are there branches occluding the light?) as well as on the conditions of other lights along the same street. Enriching the datasets collected in the field may seem to add too much extra work in the context of the specific operation at hand (in this case the simple changing of a light bulb), but the marginal added cost of gathering the few additional pieces of information could be recouped in the long run since a fuller, richer and more up-to-date picture of the city is thus made available to municipal administrators, who will therefore be in a position to make maintenance and management decisions in a more informed manner, without guesswork.

[maintenance-based updates]

Reliable

For city knowledge to be a powerful tool in urban maintenance, management and planning, the underlying datasets need to be eminently *reliable*. People need to trust that the information they are using is accurate and up-to-date. If data are perceived to be inaccurate or unreliable, they simply will not be used. Moreover, the data need to be reliably available, and not here today and gone tomorrow. The sustained existence and proven reliability of municipal datasets will create a constituency of users that will come to rely upon them for daily operations, thus perpetuating and reinforcing the need for such information.[408]

continual reliability

[data quality]

reliable sources

Depending on the source of the data, people may be inclined to rely on them with more or less confidence. What happened to me in Italy was that city administrators were loath to accept data produced by students as sufficiently dependable to be incorporated into their activities. I was forced to create my own corporation (Forma Urbis s.a.s.) to validate and integrate the student work and thus put my company's seal of approval on it, taking full responsibility for data accuracy. This process of certification was not only accepted but actually encouraged by Venetian administrators who needed a "fall guy" in case problems with the datasets were later discovered.

FORMA URBIS

[reliable or liable]

departmental certification

To expedite the whittling away at the backlog, cities may choose to enlist trusted contractors and consultants to carry out some of the work. In the long run, however, I envision a self-supporting municipal infrastructure wherein each department will be responsible for collecting and validating all of its own data, thus ensuring their reliability since the department's own

---

406 The process described mirrors my own personal metamorphosis from empirical intuition to a more theoretical reflection, which in turn parallels the evolution of planning theory itself, as described by Peter Hall in his *Cities of Tomorrow*, pp. 322 ff.

407 As in our tree projects (Creps *et al,* 2001 and Bennett *et al.*, 2001).

408 See Budić, 1994, p. 252, Table 4, under Operational Effectiveness indicators.

[implicit quality control]

operations will depend upon the accuracy of those datasets. This implicit quality control will in turn guarantee to all departments that the data they exchange with each other will be intrinsically reliable, or at least as accurate and up-to-date as possible at that time[409]. This virtuous cycle should gradually improve data reliability, accuracy and precision for the overall benefit of all municipal operations.

Flexible

rich and flexible

standard and flexible

predictable flexibility

The richness of the municipal datasets that I envision implies a degree of malleability that opens up possible avenues for reutilization of the data in disparate contexts. City knowledge is *flexible* because of its richness, and ironically also because it is fairly rigidly structured along standard coding and reference schemes. It is counterintuitive and almost paradoxical that the predictable and rigid backbone of our data structures enables the great flexibility and adaptability of city knowledge to changing circumstances. Yet the oxymoron of "predictable flexibility" is implicit in the teleological enrichment of our datasets discussed earlier. This apparent contradiction is also inherent in any situation where standards are widely adopted. The rigid abidance to the GSM telephone standard gives me the possibility to travel back and forth between Massachusetts and Italy without having to change cell phone. The universal power supply in my laptop functions in both countries regardless of whether the line voltage is 110 (US) or 220 volts (Italy). Yet life would be even easier (and more flexible) if every country strictly adhered to the same voltage and the same electrical plug configurations.

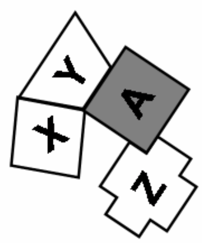

atomized flexibility

[reconfigurability]

atomized resilience

The predictably standardized nature of the data structures that encapsulate city knowledge is a prerequisite for flexible reutilization, but so is also the nature of the data that are stored in those data structures. Flexibility in our city data starts at the moment of data collection and initial data archival. Atomized data, collected and archived in the most disaggregated and fragmented – yet logical and organized – manner that is reasonably achievable will always be more adaptable and reusable than data that are aggregated, manipulated or pre-digested before storage[410]. Census data on a tract level are less flexible than those based on the block (or block group). If I want to study a specific neighborhood that straddles two tracts, I will do a better job if I can reconfigure my data using the blocks than I could possibly do by using tracts[411]. Atomized data may also offer a degree of resilience[412]. Aggregate data can often be produced even if datapoints were corrupted or missing altogether. The choice of how to make up for these flaws is left to the analyst. Aggregate data, on the other hand, frequently hide these lacunae and leave the end user no choice as to how to deal with them. It is better to obtain an imperfect set of fine grained data points than to get an aggregate

409 See for instance Craglia *et al.*, 2004 and Tulloch and Fuld, 2001.

410 "Data warehousing" is one of the pillars of the Urban Institute's National Neighborhood Indicators Partnership (NNIP). Their current attitude is to "keep the whole file at the ready so you can respond quickly as new data needs are expressed" (*Building and Operating Neighborhood Indicator Systems*, p. 36).

411 See, for instance Tufte, 1997, p. 35.

412 Resilience is one of the topics addressed by Kathi Beratan in her lecture on November 21, 2003, entitled *Managing Complexity, or the Information Needs of Adaptive Co-Management: The Durham NC air quality Case*,, in the framework of the MIT E-planning seminar, Fall 2003.

set that appears to be whole and complete. The fact is that real world data are rarely perfect.

[fine grained = adaptable]

deconstructing city knowledge

The spatial "unit of measurement" is therefore a key predictor of the flexibility of city data. In fact, the degree to which data can be converted into information depends heavily on the level of spatial disaggregation. Deconstruction of city data will always yield richer possibilities for later information building. This deconstructivist approach need not be applied only to spatial units of measurement either.

[non-spatial deconstruction]

I have struggled over the years to dissuade experts in various fields from using the prevailing "analytical" approach to data collection that entails, for instance, the evaluation of the condition of a piece of public art as a single synthetic "expert opinion" expressed symbolically as *poor, mediocre, acceptable, good, excellent*, or on a pseudo-quantitative Likert-scale from 1 to 5. The conceptual jump that experts subconsciously make when translating visual clues into their final "grading" of the state of conservation of an artifact is tantamount to gathering data at the block level but then only storing it away as a summary by tract[413]. The visible evidence that is mentally factored into the evaluation is not recorded and hence it is lost to posterity. Yet these atomic indicators could be useful for other purposes or for different analytical summations based on alternative evaluation criteria.

[visual inspection rubrics]

[atomizing is more flexible and cheaper]

Observable traces of tree disease or architectural damage are fairly easy to distinguish, as discussed in earlier chapters[414]. Non-experts, like students or staffers, can be easily trained to recognize and record these clues as long as a rubric is created to facilitate the process of visual inspection and detection. So, not only is an atomic approach to the collection of the parameters that characterize an urban element more flexible in the long run, but it is also – in the short term – less costly, since pricey expertise is not necessary during the time-consuming field work. Only after the evidence is collected by inexpensive staff, can experts be called in to conduct quality control spot checks on the validity of the collected data and also to perform more in-depth follow-ups with the objects that appear to be in the worst condition.

Reusable

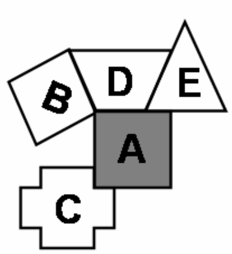

One of the main claims of the municipal knowledge infrastructure I am proposing is that it will eliminate waste and redundancy by allowing accumulated knowledge to be brought to bear as needed, time and time again, without the extra cost of collecting data repeatedly upon demand. This shift from plan-demanded data to plan-ready information is predicated upon a concerted effort to amass the necessary data using a standardized framework that allows the piecemeal addition of new records in a cumulative fashion[415]. The underpinning of this reusability lies in the rigorous structure for the labeling and referencing of data items which are spatially linked through their geographic location.

413 This represents a commingling of "facts and values" that goes against some well-respected decision-making approaches (such as Hammonds's 1980 and 1991).

414 Page 129.

415 In previous chapters, I have illustrated several examples that demonstrate the benefits of plan-ready information.

units of measurement

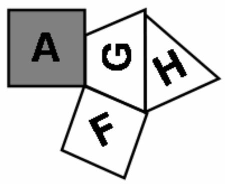

The atomization of the "units of measurement" and their geocoding onto standard basemaps creates the foundation for the reusability of the datasets at a later date, and/or under a different set of circumstances altogether. For instance, I was able to compare lagoon boat traffic data collected on two weekdays in 1997 and 1998 to determine changes from one year to the next, simply by re-using the datasets as they were and extracting just the records that were collected from the same locations in both campaigns[416]. The same traffic data are also being re-used in a higher-order project for the calibration and validation of the boat traffic model commissioned by the *Commissario al Moto Ondoso* in 2003[417].

units of analysis

Moreover, not only can the information prove its plan-readiness simply by being reusable without modification as in the examples above, but it can also be re-used and re-aggregated according to completely different "units of analysis", sometimes in ways that would have been hard to imagine *a priori*. For instance, our cargo delivery data, originally collected from individual docks, were later aggregated "by island" to determine the overall "demand for deliveries" for each isle. This lead to the plan-demanding cargo re-engineering project that has revolutionized the delivery of goods in Venice[418].

Sharable

intra-departmental sharing

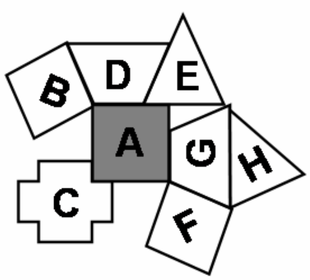

longitudinal legacy sharing

Re-usability, richness, flexibility and reliability, empower us to share information with others. First and foremost, the data could be shared internally <u>within</u> a division or department. Even such a straightforward form of sharing is not quite the norm today, despite its apparent simplicity. For instance, a WPI team proposed to institute a simple intra-departmental form of communication between two divisions within the Traffic and Parking Department in Cambridge, which saved the city upwards of $300,000 a year that would have been lost in meter jams[419]. Obviously, such patently advantageous forms of sharing are the easiest to put in place, since they instantly pay for themselves. Without such venal incentives, WPI teams practiced a form of intra-departmental sharing all along, both in Venice and in Boston, by exchanging and bequeathing our internal legacy databases from one team to the next over the years, parlaying previous successes into stepping stones that would catapult new projects into more ambitious undertakings. Thus we were able to gradually and systematically build upon past projects toward bigger and better results that are truly changing the way things are done in Venice. Thanks to our internal sharing, we can truly claim that we are "leaving Venice better than we found it", and we are well on our way to do the same in the greater Boston metropolitan area and in Worcester as well.

416 Carrera 1997, 1999.

417 *Forma Urbis* was contracted by the *Consorzio Venezia Ricerche* to create the informational foundation for the model. Phase one was completed in March of 2004.

418 as discussed in detail on page 37 and page 118. Our plan was featured in the September 27 issue of *New Scientist* and was also the subject of my interview with the BBC World Service Radio on October 2, 2003.

419 See Footnote number 361.

[the framework for sharing]

We also shared our outcomes and methods with numerous agencies and organizations in Venice and in the US who are all benefiting from our work. More importantly we have created the premises for further sharing across agencies by, for example, creating a *de facto* standard for the coding and labeling of Venetian canals. Insula S.p.A. and various departments in the city of Venice, as well as offices in the Provincial and Regional government have all shared information in an effortless manner thanks to the standard reference system for canal nomenclature that I developed in the 1990's. Similar forms of sharing are occurring in many other areas as well.

inter-departmental sharing

Once the municipal information framework is created, sharing is possible among different departments as well. Such arrangements could first of all follow "paths of least resistance" by concentrating on operations where inter-departmental sharing is already a reality due to institutional mandates as is the case between the Boston Inspectional Services Department and the Historic Districts Commission whenever a building that requested a particular permit is a registered historic property[420]. This is an example of institutional sharing that is already in place and is mandated by law wherein City Knowledge solutions could be easily introduced with instant benefits and without revolutionizing standard practices[421]. After the "low hanging fruits" have been addressed, sharing could be treated as an additional instrument at our disposal to maximize operational efficiency. In an advanced sharing framework, common resources could be mainstreamed into municipal operations in such a way as to exploit synergies that are completely untapped today due to the disconnects that exist among departments.

[paths of least resistance]

[low hanging fruits]

sharing beyond the firewall

Beyond inter-departmental sharing, one can foresee the possibility of making some of the information available to interested parties outside of the municipal firewall. Some information could, and possibly should, be made available to citizen groups, both in its raw original formats and in pre-digested versions for public consumption, as is done in the flourishing community statistical efforts like "Neighborhood Knowledge" in California[422] and the National Neighborhood Indicators Partnership (NNIP)[423]. We experimented with this type of sharing with the citizens of the island of Pellestrina (Venice)[424] and, through the Environmental Protection Agency (EPA), with the Chelsea Creek communities of East Boston and Chelsea, Massachusetts[425]. Beyond public constituencies, the other major users of municipal data are probably academics, especially in the fields of public policy and urban studies and planning, who are always scouting around for information to support their research interests, as I personally continue to do to this day.

[neighborhood data]

[sharing with academia]

emergent system

Sharing will allow the distributed data producers in all of the municipal branches to connect together to form a virtual intelligence that is

420 Hart *et al.,* 2004.
421 *Idem.*
422 www.nkca.ucla.edu
423 www.urban.org/nnip/
424 Battocchi *et al.*, 2003.
425 Desmond *et al.*, 2002.

greater than the sum of the parts and that could surface progressively and spontaneously in whatever way is most suitable at any particular time, gradually taking on the connotations of a true *emergent* system with pleasantly unpredictable, yet promising prospects for improved municipal services that will benefit individuals as well as government and non-government organizations.

Secure

The possibility and desirability of allowing departments to share information with each other brings up the specter of the information falling into the wrong hands. Having a detailed layout of all of the city's gas mains could greatly facilitate the work of road crews and minimize possible "dig safe" hazards, but the same map can also become a weapon in the hands of local disgruntled Americans or of foreign disgruntled anti-Americans alike. In some ways, true homeland security demands the complete knowledge infrastructure that I am proposing, but at the same time all this knowledge could also become a bane if misused by terrorists or criminals. In my view, ignorance is not an option. So instead of refusing to embrace City Knowledge because it could be a dangerous tool in the hands of al-Qaeda, we should just make sure that adequate safety precautions are taken to ensure the security of the data cannot be compromised.

[dig-safe]

[homeland security]

[passwords, certificates and authentications]

As is commonplace in ICT systems, sharing of City Knowledge will be controlled by a system of passwords, certificates and authentications, and the same safeguards that protect the thousands of other sensitive internet sites should be adopted to protect vulnerable urban information repositories. Varying degrees of protection can be placed upon municipal data, depending on its nature, so as not to unnecessarily burden the system with too much security when it is not warranted. When it is appropriate, the data should be protected in the most suitable a cost-effective way. If adequate defenses cannot be put up, it may be better to remove the data from any system that is accessible through internet connections until a safer system can be put in place.

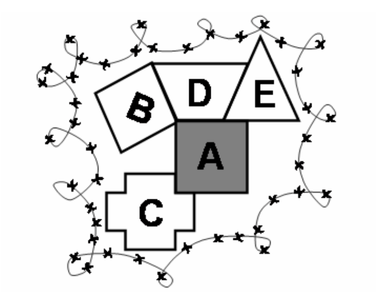

knowledge and safety

In the end, though, the possible costs associated with protecting city knowledge should not keep us from developing that knowledge in the first place. In fact, it is my hope that the growing concern for the safety and security of urban populations will spur a well-funded effort to embrace city-knowledge principles all around the world[426]. If a natural disaster or criminal incident should occur, emergency crews will want to know exactly everything there is to know about the location where the crisis is occurring. For example, we helped the Boston Fire Department keep track of all Underground Storage Tanks (USTs)[427] and all tall buildings[428] so that fire crews could know what to expect if an emergency was called at a site where either a UST or a tall building was located. In case of evacuations, for instance, fire crews would like to have an approximate estimate of how many people are expected to be in a particular building, and they would also want

426 Witness our involvement with the Italian *Protezione Civile*, as mentioned in footnote n. 275.

427 O'Donnell *et al.*, 2002.

428 Gaewsky *et al.*, 2003.

to know whether and where hazardous materials may be stored within that building, so that the proper countermeasures can be taken.

privacy

Safe and *secure* city knowledge is a prerequisite for the system to work smoothly even in the absence of any terrorist or malicious threat. Simple protection of privacy, for instance, mandates that safeguards be put in place to prevent the release of personal information. All in all, a gradual transition from stand-alone systems, disconnected from the internet, to intranets with impassable firewalls (to share information internally), to limited internet accessibility with proper certification and authentication (to let selected outsiders get access to your information) will smoothly get each department's portion of the overall municipal information system to a proper level of security without risking violations of privacy or malevolent attacks.

[gradual development path]

onward

Attainment of the qualities discussed in this chapter will help ensure that our municipal knowledge infrastructure does not become another exercise in futility, but it actually stands a real chance of becoming an irreplaceable tool for urban maintenance, management and planning. Most municipal efforts will experience difficulties in one or another of the areas highlighted above, but the important thing is to set the system in motion and allow emergence to take place gradually at its own pace. The following chapters provide additional guidance on how to steer the process into a fruitful direction.

FOUNDATIONS OF CITY KNOWLEDGE

information-conscious modus operandi

The qualities discussed in the preceding chapter define the dimensions through which we can evaluate the performance of the knowledge infrastructure that I am proposing. Varying degrees of success can be expected in the pursuit of each of those desirable characteristics of a City Knowledge system, but we have identified practical mechanisms that can at least facilitate the attainment of such positive traits. Once the primary paradigm shift has taken place toward an information-conscious *modus operandi* in all municipal activities, City Knowledge can begin to emerge, based on the following foundations:

- Distinct informational jurisdictions
- Distributed, atomic data acquisition and organization
- Sustainable update mechanisms
- Institutional and/or voluntary sharing of information
- Interagency coordination

"glued" together by a

- Middle-out approach to the development of City Knowledge

[gradual building of foundations]

In the sections that follow, these fundamental principles are each described in detail. Although these elements must all be in place in order to extract the totality of the benefits that City Knowledge can offer, they can also be attained partially and sequentially in a gradual progression toward the ultimate goal of a full-fledged City Knowledge system[429].

[no City Knowledge systems in existence]

As the most recent literature indicates[430], such a comprehensive system does not seem to exist in its entirety anywhere in the world, and even where I was personally engaged with the inner workings of municipal departments – namely in Venice and in various Massachusetts cities as described in parts II and III – these principles have not been completely adopted – not even unofficially. This chapter is therefore purely conjectural, although it pulls together pieces of the puzzle gathered from each of the various cases presented earlier, as well as from comparable efforts discussed in the urban studies and planning literature.

[conjectural hypotheses]

[obstacles]

Numerous obstacles will impede the emergence of a full-fledged City Knowledge system, as described earlier[431]. There is ample literature in the field of Management Information Systems (MIS) that describes the difficulties in the diffusion of information systems in organizations[432]. There is also a growing body of academic literature specifically dedicated to the

429 The sequencing is somewhat flexible. Some forward-looking municipalities may start by setting standards. Most practical-minded towns have generally started some sort of computerized bottom-up data collection and GIS layers (ICMA survey, 2002).

430 Laurini, 2001; Brail and Klosterman, 2001; Geertman and Stillwell, 2004.

431 Starting on page 165.

432 For example: Laudon and Laudon, 1996; Marchewka, 2003, pp. 4-5;

institutionalization of GIS in organizations[433]. Even after the adoption of such an approach has been agreed upon, the devil is always in the implementation details. It is still unclear "how the management in a public agency can move the organization/jurisdiction through the process of […] GIS implementation"[434]. The consensus opinion indicts the "neglect of the human and organizational aspects" as the primary culprit in the failures of the past[435]. In the municipal domain where my dissertation is focused, the "technology-push" from the top has failed to penetrate deeply enough into the frontline offices where fine-grained data are the daily currency of operational decisions for maintenance and management activities. My middle-out strategy proposes to switch instead to a "demand-pull" approach[436], which centers on the needs of the periphery first.

[where City Knowledge will not work]

City Knowledge is not a panacea and it may be better-suited for some municipal governments than for others. Small towns may not have the resources to even consider this approach[437]. Big city government may be too fractured to be conducive to it[438]. Since space plays a key role in pulling together my City Knowledge strategy, there is an implicit reliance on technology (GIS and RDBMS in particular), which may be utterly inappropriate in some contexts, such as poverty-stricken locales, where basic survival needs overwhelm any other municipal service. City Knowledge may be a luxury that only communities in the more affluent parts of the world can really afford[439]. In some other situations, the concept of a comprehensive municipal information system may be deemed undesirable for political or cultural reasons. With these disclaimers in mind, in the sections that follow I focus on how City Knowledge could be embraced as a strategy and a *modus operandi* applicable only wherever the circumstances permit this concept to be viable and desirable.

[pioneering department]

In this dissertation[440], I submit that it may be possible for one municipal department to set the process in motion by single-handedly beginning to systematically collect data that unequivocally fall under its

433 For example: Budić, 1994; Campbell and Masser, 1995; Reeve and Petch, 1999; Azad (1998) provides a thorough review of the process of managing GIS implementation.

434 Azad, 1998, p. 18.

435 Reeve and Petch, 1999, p. 5. Evans and Ferreira (1995) however suggest that technology is also culpable.

436 *Idem.*

437 Although I have begun to develop a project proposal to connect small towns (like my hometown of Spencer, MA) with Regional Planning Authorities (like the Central Massachusetts Regional Planning Commission – CMRPC) through web-based applications created and supported on the RPA's central server that will automate routine tasks such as construction permitting at town hall while "informating" planning at the regional level.

438 Although, when taken one department at a time, even the biggest bureaucracy can be tamed, as my cases showed. The biggest obstacle may come when dealing with standardization and jurisdictional coordination, which have not been institutionalized explicitly even where I worked directly with willing departments.

439 Yet, if we consider information as an infrastructure as I am proposing, we may want to put in place information systems at the same time as we lay down other infrastructural essentials, such as water, electricity and sewers. If nothing else, we could at least keep track of these physical systems more efficiently, thus minimizing maintenance costs in the long run.

440 Grounded on some recent literature surrounding the "tipping point" principle (Gladwell, 2000 and Godin, 2001).

[capture backlog]

domain of control until it has organized its own backlog of existing information in an exhaustive and systematic manner[441]. Having caught up with the pre-existing *status quo*, this pioneering department can then begin to implement procedures and invent mechanisms through which any new pertinent information that comes into being can be captured as soon as possible and with the smallest expenditure of human and financial resources.

[intercept change]

spreading city knowledge

After setting the example by creating a working operational model of how to manage its own slice of city knowledge, this innovative department could then set in motion a chain reaction[442] that could "infect" other departments[443]. Whether this chain reaction happens or not, my claim is that the mere systematic organization of a single aspect of a department's duties would justify the effort in and of itself, so any additional benefit would be a value-added "cherry on the cake".

[spontaneous coordination]

Beyond that, if two departments could demonstrate the viability of a full-fledged departmental City Knowledge system, the need or desire to coordinate a bare minimum of citywide standards would emerge naturally. Such a harmonization would in turn ensure that, as the various departmental systems come on line, they can begin to share information on a "need to" basis, starting from the existing mandatory bilateral collaborations that already bind them into institutional relations with each other.

[institutional sharing]

[full-fledged emergent system]

Over time, this gradual growth of city knowledge from one department to the next could eventually encompass the entire municipality and enlist government offices, civic associations, academic institutions, non-government organizations and even private citizens and businesses in the continuous upkeep of urban information[444].

the rest of this chapter

The following sections expand on these fundamental processes that in my view will foster the emergence of sustainable City Knowledge. At the end of each of the six sections that follow, I critically assess how my particular approach – culled from specific personal experiences I had in Venice or in Massachusetts – compares with past efforts of similar nature. I try to identify what makes my approach different and/or better than alternative methods and try to distill what aspects of current practice would have to change (and how) in order to move a municipality in the direction I propose.

[empirical basis]

[similar to other efforts?]

[different/better?]

[change what and how?]

My numerous experiences in real world implementations have convinced me that truly useful, complete and reliable urban information systems are indeed within reach of a typical municipal government in a medium-sized, "western" city like Venice, Cambridge or Worcester. Often,

441 This process may be achieved with or without the aid of an MIS department. Most medium to small towns will never have the benefit of such a department, yet I have personally experimented with several alternative ways for such small towns to embark in the systematic accumulation of city knowledge. See footnote 437.

442 Like Malcom Gladwell's *Tipping Point* (2000).

443 Like Seth Godin's *ideavirus* (2001).

444 Of course, there is a pretty huge leap between simple bilateral sharing and citywide sharing. Coordinating these efforts gets progressively hard as n>2. My argument is that ifbottom-fed middle-out initiatives can overcome the complex hurdles of multiagency coordination then the resulting approach will be necessarily long-lasting and resilient.

[MIS and GIS theories]
[diffusion and implementation]
[individual aspect is proven]
[ensemble not proven]
[the middle-out recipe]

what my experience has led me to propose as a foundation of City Knowledge, coincides fairly closely with existing concepts in the area of MIS diffusion, or to parts of other approaches in the GIS implementation literature. As one would expect, not a lot of what I propose is completely novel. On the other hand, my personal experiences in the field – as well as those published by many fellow urban researchers – have at least revealed the merits of each individual foundation of City Knowledge.

What I cannot prove, but I instead predict – based on the evidence presented herein – is that these five ingredients, selected and combined in a fairly flexible manner, can produce results that far surpass what you would get by simply mixing the constituent parts. The novelty of my overall proposal is in the nuances of each ingredient and in how the ingredients come together into the final product.

Together, the foundations that are described in the rest of this chapter can elevate municipal data to urban information and thus contribute to the sort of city knowledge that would support second-order analyses and complex decisions. But first we need to introduce the technique that makes it possible to put all these ingredients together: the *middle-out* approach. This method is not exactly like a culinary recipe from a cookbook since no specific sequence of steps and carefully measured dosages are required to put the five ingredients together.

The result is also not guaranteed since – as far as I know – nobody has quite put all of these ingredients together in the same concoction quite yet.

The "Middle-out" Approach

The "middle-out" approach to City Knowledge combines the self-serving ingenuity and energy of bottom-up initiatives and the coordination and standardization of top-down approaches. Middle-out is the methodological framework which allows City Knowledge systems to emerge. It is a consciously unobtrusive strategy which guides the choice of common-sense, low-impact tactics that gradually translate the promising concept of City Knowledge into an achievable reality.

bottom-up enlightened self-interest

This hybrid approach empowers the more peripheral branches of the municipal hierarchy and gives front-line civil servants a stake over the upkeep of the information that they themselves require. The "middle-out" approach assumes that municipal officials are motivated by a desire to improve their own performance at their daily tasks related to the maintenance of city property, the management of city services, or the planning of city futures. The bottom-up component of *middle-out* exploits the all-too-human self-interest that civil servants undoubtedly harbor, like the rest of us. Everyone wants to do a good job and receive praise and recognition from peers and superiors. Enlightened self-interest will be the motor that will feed the information infrastructure that I am describing.

top-down teleological coordination

To rein in the anarchy that might ensue if each civil servant with a modicum of computer savvy were to create an information system from the ground up – as has been somewhat the case in the last decade of the XX century after the advent of personal computers and the mainstreaming of databases and geographic information systems[445] – the *middle-out* approach entails a degree of top-down coordination and management. Information systems will be indeed developed in bottom-up fashion by the front-line departments where the action really happens and where information is produced and consumed on a daily basis, but these efforts will be managed and coordinated at the departmental or divisional level and will include a teleological (forward-looking) approach to avoid unbridled duplication of effort and overlapping of jurisdictions that will hamper the sustainability of the City Knowledge enterprise in the long run.

critique of top-down and bottom-up

[shortcomings of top-down]

Both top-down and bottom-up are fundamentally unsustainable in their pure incarnations and both frequently lead to a waste of energy, time and money[446]. Pure top-down approaches to the diffusion of information systems in organizations are generally poorly received because they fail to engage the rank and file[447]. After huge initial investments, they struggle to achieve the fully-integrated coordination and seamless operation that they promised, especially when the task is not easily automated and the organizational structure is not very hierarchical[448]. Nevertheless, the military is a glaring example of where such a hierarchical approach could indeed work – possibly more reliably than the emergent middle-out system I propose.

---

445 Geertman and Stillwell, 2004; Brail and Klosterman, 2001.
446 Reeve and Petch, 1999, pp. 4-9.
447 *Idem.*
448 See Marchewka, 2003, especially pp. 4-6; Reeve and Petch, 1999, p. 76.

[shortcomings of bottom-up]

Grass-roots, bottom-up schemes work well for a while, as long as some "champions" keep the efforts going, but the resulting information systems eventually fall by the wayside because they refuse to "grow up" or are unable to connect with other systems, or simply fail to keep up with technological or organizational changes[449]. Bottom-up efforts frequently fail because of lack of sustained support by a dedicated core of users/developers and because of the related lack of adequate funds to keep the effort alive through changes in personnel, equipment, and software tools[450].

a pragmatic hybrid compromise

The middle-out method is predicated on striking a careful balance between the potentially oppressive and unwelcome rigor of a typical top-down hierarchical system and the practical ineffectiveness of isolated bottom-up efforts, which almost inevitably fail to scale up or to integrate with other similar systems[451]. It represents a pragmatic compromise that taps into the positive aspects of the two competing approaches, leveraging the energy and self-interest that power bottom-up endeavors, and endowing them with the structure, reliability and sustained financial support that accompanies top-down enterprises.

pseudo- and proto- middle-out

Some development paths being recommended today resemble this middle-out approach[452]. As mentioned[453], the City of Cambridge, Massachusetts has begun an effort of this sort, by distributing GIS specialists in the main city departments under the orchestration of the MIS department, although that endeavor is still in its infancy. This is one of many possible paths to promote the diffusion of GIS in municipal operations. The Management Information Systems (MIS) literature has a lot to teach us about the variety of methods that have been tested to favor the dissemination of information systems in organizations[454]. My middle-out approach would be more similar to the Digital Earth effort[455], but limited, at first, to individual municipal boundaries and to the agencies operating therein. Ferreira, in particular, has championed a variation of this middle-out approach when he proposed the use of lookup tables to correct on-the-fly the "stubborn" standardization errors that regularly appear in municipal datasets[456], as well as (with Evans) when he discussed a more general approach to the "messy" technical and organizational issues confronting GIS today[457].

---

449 Klosterman, 2001 and Singh, 2004.

450 *Idem.*

451 See the interesting "Problems in the Middle Layer" that Reeve and Petch (1999) discuss on page 25 of their book on GIS and organizations. My overall City Knowledge approach is intended to make the "theory" illustrated in figure (a) more achievable, to overcome the "reality" in figure (b).

452 See for instance Barr, 1991; Campbell, 1999 and Yeh, 1999; Talen, 1999; Ferreira, 1998 and especially Reeve and Petch, 1999, p. 50.

453 Page 125.

454 For instance Laudon and Laudon, 1996; Avison and Fitzgerald, 1992; Flynn, 1992; Reeve and Petch (1999) argue that "the development of GIS is like watching a video of the history of conventional information systems being replayed at fast forward speed" (p. 1).

455 See http://www.digitalearth.gov/ and Crockett, 1998.

456 Ferreira, 1998.

457 Evans and Ferreira, 1995.

What these new methodologies have in common is a recognition that "GIS technologies are *not* divorced from the interplay of organizational life: rather they are subject to its vagaries and power relationships"[458]. A middle-out approach should not only simplify the more technical pitfalls of pure top-down and bottom-up approaches[459], but also promises a more gradual, hence smoother, and less traumatic path for the organizational transformations needed to ensure a widespread acceptance and a successful adoption of the principles of City Knowledge in municipal agencies.

horizontalizing growth from the middle

The hybrid combination of these two extant methodologies, not only exploits the good traits of both, leaving the negative connotations behind, but it also removes the hierarchical verticality of the previous constructs, which was in itself a source of tension between management and rank and file across the entire municipality. The locus of control of information is now at the level of a division or department[460]. Within each department, managers and staff are partners in the development of information systems that improve that department's effectiveness, with no imposition from the mayor's office or the city manager or the MIS department[461]. With a proper consensus-building approach, the entire department or division can bond together around the common goal of producing effective computer-aided tools to streamline every function that the department is responsible for[462].

["slices" of City Knowledge]

[City Knowledge as a "compendium"]

The *esprit de corps* that will result from this devolved information design will give staying power to each department-based "slice" of city knowledge. These self-contained units could then connect horizontally with other departments or divisions across the municipality. Top-down coordination among all municipal branches and bilateral agreements between individual departments could thus enable the overall city knowledge "compendium" to emerge from the middle out as a combination of these self-directed, semi-autonomous bottom-up efforts[463].

the tenets of middle-out

Recapitulating, the basic tenets of the middle-out approach are:

1. front-line offices are directly vested into the collection and upkeep of their own city data, information and knowledge[464];
2. each office will systematically address each of the city elements over which it has birthrights and assess the costs and benefits of creating a comprehensive knowledgebase relating to each element;
3. once the low-hanging fruits are identified, the office will explore all feasible and reliable means for the accrual and sustainable upkeep of the data, starting from no-cost options;

---

458 Campbell, A. J.. 1999, pp. 621-631.

459 Keating *et al.*, 2003.

460 Which is where the reward structures would have to also be located (see Singh, 2004).

461 Although the department may be assisted by MIS in the development and implementation.

462 Some of these issues are covered in Nedović-Budić and Pinto, 1999; Campbell and Masser, 1995; Reeve and Petch, 1999, ch. 6.

463 Cf. Barr, 1991.

464 Though they may yet delegate the nitty-gritty of the data manipulation to a technical department as was done for instance in Calgary, with its "Data Utility" concept (Findlay, 2002).

4. data are collected at the finest appropriate grain and the same office is also responsible for the efficient and effective upkeep of the data;
5. the front line office participates in intra- and inter-departmental caucuses for the definition of top-level issues, such as standardization, jurisdictions and sharing;

what's new about middle-out

Middle-out is a strategic approach that relies on opportunistic tactical choices by self-interested front-line users (the "ants") who engage in self-organizing behavior to take care of their impellent needs. The middle-out methodology is more focused in the implementation of a comprehensive municipal information system than in the diffusion of the *information-aware modus operandi* that subsumes the City Knowledge philosophy.

[empirical basis]

All of my examples from Italy and the U.S. were carried out in collaboration with specific departments or agencies. We purposely shunned projects imposed from top management or politicians and we did not work with individuals except in the case of one-person institutions[465]. All arrangements were conducted with department heads or program directors. Our projects were institutionalized at the department level.

We were able to show some evidence of horizontal diffusion thanks to shared reference codes both in Venice[466] and in Cambridge[467]. We have also shown how our approach can spread horizontally across department or even municipal boundaries, as evidenced by our ability to elicit interesting and challenging projects from new sponsors by showing them what we did elsewhere and how our sponsors benefited from our studies. To some degree, we expect that the desire to embark in a systematic inventory of physical assets under a department's jurisdiction ought to emerge spontaneously in the various branches of municipal government. In any case, I posit that the diffusion of this approach could be abetted by some lobbying by "communicative planners" who have a lot to gain from the implementation of such systems. Yet the primary focus of the middle-out approach remains implementation and not diffusion[468].

[similar to other efforts?]

As amply discussed in the sections above, there have been several efforts in the recent past that contain seeds of my middle-out approach[469]. To date, none of the cities – or even departments – with which I collaborated have formally espoused my wholesale middle-out approach. In some sense, middle-out is one of the more conjectural aspects of my thesis. I am proposing that it would be beneficial for a municipal department to

465 Such as the Boston Air Pollution Control Commission (APCC) or the Boston Conservation Commission (BCC).

466 As a consequence of the canal UNESCO-MURST project (Carrera, 2001b); the transferal of the base layers to Insula (Carrera, 1999d); the traffic projects (Carrera 1996, 1999a); and of some professional follow-ups, such as *EasyDocks* and *EasyBridge* applications mentioned earlier.

467 With the initial cataloguing of parking meters in the Traffic and Parking department (Cullen *et al.*, 2002) that spread to both the meter maintenance crew and the parking control officers (Flynn *et al.*, 2003);

468 Although I find the Tipping Point (Gladwell, 2000) and Ideavirus (Godin, 2001) concepts promising agents of diffusion.

469 See footnote 452 and related sections.

consciously adopt my approach, yet I personally have no evidence that it would indeed work[470].

[different/better?]

What is new about middle-out today is that recent developments in information and communication technologies – especially spatial tools – now enable this approach to be cost-effective at the department or division level, which gives the beneficiaries direct control over the tools used to represent reality[471]. Physical objects and dynamic activities can therefore be captured and stored in their geographic location by the agency most directly engaged with them[472]. Moreover, GIS and databases have been around long enough that the need for some top-down coordination is more commonly acknowledged.

Perhaps what is novel about my approach is simply that I am introducing a new metaphor to empower municipalities to consciously foster emergent behaviors so they materialize from the front-lines out. I am packaging my concepts in new re-descriptions that I hope will strike other planning practitioners as inevitable[473] and will thus become vehicles for the creation of the sort of City Knowledge systems that my numerous examples collectively suggest.

[change what and how?]

The path to adoption of middle-out starts with the conscious decision to apply this method in the framework of the department's information strategy. The issue here is not the simple "adopt/do not adopt" choice, but a more authentic commitment to City Knowledge and its principles – described in the sections that follow – that will guide the success of the middle-out approach[474].

[470] Although others do find some positive evidence using approaches that are similar to mine. See for example Keating *et al.,* 2003; Innes and Simpson (1993); and Reeve and Petch, 1999.

[471] A desirable characteristic according to all recent implementation theories as described in Azad, 1998, Reeve and Petch, 1999, Nedović-Budić and Pinto, 1999, Singh, 2004.

[472] Although typical implementation paths may still be adopted, such as those in Azad, 1998 and Reeve and Petch, 1999, with the technical support of the city's MIS department, the assistance of outside consultants or the intervention of regional planning authorities.

[473] See "Intellectual Contribution" on page 15.

[474] More on this starting on page 230.

Informational Jurisdictions

jurisdiction over structures and activities

A very fundamental tenet of City Knowledge is the attribution of clear "informational jurisdictions" over the "birth and death" of the "structures" that make up our urban environment, as well as over the "actions" that occur within it[475]. One should not confuse the concept of informational jurisdiction with some form of *exclusive* responsibility over the maintenance and management of these city elements by a single entity or department[476]. Nor, should our informational jurisdictions be mistaken for an assignment of complete "ownership" over all information that pertains to a specific item[477].

birth and death certification

[parental authority]

In the domain of physical structures, the definition of clear jurisdictions that we propose refers solely to the "parental authority" over the certification and formalization of the birth and death of these objects within the municipal administration.

[administrative existence]

[bureaucratic acts accompany actions]

The notion here is that the physical assets that make up the "public" part of a city – roads, buildings, trees, parks – have an "administrative existence" that begins and ends with official bureaucratic acts that always accompany operative government actions in the real world. So, for instance, the parks department normally puts out a work order for the planting of a tree somewhere in the city. The administrative red-tape (the work order) is converted into an action in the real world by a crew that physically plants the tree into the ground. The bureaucratic birth of a tree is marked by a work order (or by a contract if the service is outsourced) and the tree's physical appearance into the real world is a direct consequence of the fulfillment of the work order or contract[478]. In order to keep track of this tree for the rest of its existence, we will use a standardized informational scaffolding to record the tree's characteristics, as explained earlier[479]. The administrative birth of a tree could thus be treated in a manner that is not too dissimilar to how our birth is also "tracked" by an official administrative act – namely our birth certificate – when we are physically born into the world[480].

naming and coding

An arborist's recordkeeping starts simply with the assignment of a label or code to each tree to identify it unambiguously vis-à-vis all other trees in the city, just like our life starts by our receiving our names from our parents and by our being tagged with a bracelet so as not to be "switched at birth" in the nursery room.

birthrights

In the case of the "urban forest", I am hypothesizing that the parks department might have "informational birthrights" over trees[481] and that it would thus be its responsibility to maintain the corresponding GIS layer,

475 Akin to what Nedović-Budić and Pinto (1999) call "Responsibility", page 58.

476 *Idem.*

477 *Idem.*, see "Ownership".

478 Obviously this only applies to "public" trees. A canopy study like the one described starting on page 125 would thus still need to capture the canopy of "private" trees as well as those of trees in a group like in a forest (not individual trees), which may not be recorded individually at least initially.

479 See the Cambridge example at page 125 and following.

480 We have all heard the anecdotes about "ghost" people who were not considered "alive", despite their obvious physical existence, because they could not exhibit a valid birth certificate.

481 This is merely an hypothesis used for the sake of the foregoing discussion. It is irrelevant which specific department would really be in charge of trees in a specific city.

since it knows exactly where a new tree is planted and it is also in a position to give new trees a unique name or code to identify them from that moment on. The parks department may also want to attach a few other permanent pieces of information to this "tree birth certificate", like the species of the tree and the date of planting. These birth records are a necessary foundation to city knowledge although they are not sufficient, by themselves, to support all of the desirable re-use capabilities that were discussed in previous chapters.

[birth certificate]

deathrights

Despite the longevity of many plant species, even trees do die, either due to natural causes (old age and disease) or due to sudden traumatic events (lightning, tornadoes or violent impacts with trucks or cars). Not infrequently, trees are also purposely cut down for one reason or another. Some department needs to be given "deathrights" over the recordkeeping related to the removal of the tree from reality. Just as we have "death certificates" that formalize our passing from this life, some department needs to have the authority to remove the tree symbol from the tree layer in the municipal GIS, so that everyone will know that the tree no longer exists.

[death certificates]

The assignment of very clear jurisdictions over birthrights and deathrights for each element of the municipal infrastructure is a *conditio sine qua non* for this whole system to function properly over the long run. It is possible that quibbles may arise among departments concerning the assignment of jurisdictions over birth certificates, but the most clear-cut way to determine them would be to look at who in actuality is responsible for installing, creating, constructing, or giving permission to erect or establish each category of physical objects. Conversely, it may be that the jurisdiction over the certification of the "death" of that object may reside with the same department, or with whichever other department is in charge of physically removing, dismantling, eradicating, uprooting, destroying and/or disposing of the object.

[jurisdiction over birth]

[jurisdiction over death]

[parent department]

In the end, only one parent department or entity should emerge clearly as the responsible party for tracking the birth of each of the physical assets that make up a city. Similarly, the same or another department will be in charge of certifying the administrative death of objects in a specific class, thus allowing the city to track the entire life cycle of physical assets owned, operated, maintained or managed by the city[482].

[life-cycle tracking]

jurisdiction over activities

Jurisdictions are really important for city structures, because they entail the potential for keeping track of changes after the backlog of pre-existing assets has been tackled. When it comes to "activities" (like traffic, parking, crime, or economic development), jurisdictions relate more to who manages or regulates the activities, so birthrights and deathrights apply only marginally here and the assignment of responsibility for the tracking of changes to these activities relates more to the management of periodic updates than to the tracking of individual birth or death events. Some events have actual births (which fall on the day that they occur). Events also have surrogate parents, such as "masters" of ceremonies[483], parade organizers, and

482 The so-called GASB-34 accounting guidelines will make such recordkeeping more and more commonplace for financial purposes (http://www.gasb.org).

483 In Italy we call the main organizer of an event the *padrino* or *madrina* (depending on the gender).

the like. In these contexts it is possible therefore to attribute the responsibility *in loco parentis* to some individual or agency. Other dynamic phenomena are unpredictable and fickle, so they need to be adopted by some agency that is put in charge of monitoring and managing the activity, from traffic to land use, to economic development.

[relational linking]

From the moment the tree is born into the municipality (e.g. from the moment it appears on a shared web-GIS system), until its "official death", different departments will be able to attach assorted pieces of information to this tree, by creating relational databases that can link to the tree through its unique code assigned at birth. The ability to link up information to an object may depend on permissions granted by those who have jurisdiction over it. Overall accords about "read" and "modify" rights might also be agreed upon by a committee that coordinates the creation of the City Knowledge system[484]. In general, linking ought to be universally permissible to anyone who wants to attach his/her datasets to publicly-owned properties.

[open-code doctrine]

What we are proposing here is that everybody ought to call each object by its "official name" and, at the bare minimum, all of the reference codes and perhaps the corresponding map layers should be completely in the public domain. This "open-code[485]" (and open-layer) doctrine will foster the independent development of several proprietary datasets developed by individuals, academics, NGOs, private companies as well as other government agencies, with the major benefit of creating an underlying capability for sharing that does not exist today. Nobody will have to invent new names for these objects, nor will there be any need for people outside of the municipality to create their own GIS maps of the objects[486].

[read, write and modify permissions]

[“field-level” permissions]

Beyond the sharing of mere codes and spatial locations, selected portions of municipal datasets may also be made directly accessible to outsiders, through a system of permissions. Authorized users could be granted "field-level" permissions that would specify exactly which fields in which database are accessible to whom, for reading, writing or modifying[487].

what's new about these jurisdictions

The originality of the concept of "informational jurisdictions" that I am proposing is not so much in the insistence on the definition of boundaries between departments, which already exist and may actually be a potential detriment to the achievement of City Knowledge. The concept is more subtle than that. The modifier here is key. The "informational" borders around elements of our urban landscape that I am proposing are

---

484 See section on "Overarching Standards" starting at page 219.

485 The concept of "open code" is akin to that of "open source", where programming code is made transparent to users who are thus encouraged to modify it and in turn share it with others in a never-ending virtuous cycle of successive improvements. The Linux operating system is probably the best known of the open-source applications circulating in cyberspace. In this paper, the phrase "open code" actually refers to the sharing of reference codes, i.e. nametags, that uniquely label real-world objects subject to municipal maintenance, management or planning.

486 The Ordnance Survey in the UK is a leader in the development and distribution of "mastermaps" that provide the entire country with uniform base maps and unique topographic identifiers (TOIDs) as a universal foundation for the sharing of spatial information (http://www.ordnancesurvey.gov.uk). More on TOIDs on page 222.

487 See section on "Information Sharing" starting at page 214.

more specific – more spatial – than the generic departmental boundaries that already exist. In some towns all information may be considered the purview of the MIS department, if there is one. In smaller towns, information may simply not be on the agenda at all, so that all transactions with the public may be limited to acquiring "documentation" to support an impending decision. I propose to delegate the informational responsibilities to front line offices that oversee the "birth" of the element if possible.

[empirical basis]

We proposed this information-centered approach to defining jurisdictions in Venice as well as in Massachusetts. In our experience, though, it has been applied not as a "top-down" agreement among agencies defining the formal areas under each agency's control. Despite the lack of coordination from the top, the intrinsic nature of departments creates *de-facto* spheres of influence that generally match the informational jurisdictions as defined herein. Some of the clearest jurisdictional boundaries in Venice are those of Insula, that was created specifically with the charter to maintain the canals and infrastructures of the city[488] and the *Commissario al Moto Ondoso* who was appointed specifically to overcome traditional jurisdictional overlaps in dealing with the traffic problem in the canals and lagoon of Venice.

In the U.S., we really focused on jurisdictions with a recent project which explored the application of City Knowledge principles in the Environment Department of the City of Boston[489]. By analyzing the information flows in the key processes of the Boston Conservation Commission and Boston Landmarks Commission, the team was able to propose streamlined on-line procedures for the management of permit applications. Although this system has not been implemented yet, it is one of the better examples of a real-world application of this and other City Knowledge principles[490].

[similar to other efforts?]

Although several scholars have recently discussed similar concepts[491], the primary novelty of our approach lies simply in the renewed focus on information and space. In this *information aware* context, we can then be more overt about assigning jurisdictions over the accrual and updating of the data.

[different/better?]

In the simplest situations, my proposal gives primacy to the spatial aspects and therefore assigns the jurisdiction to the front-line offices that are actually interacting with the physical world and are doing the hands-on installation or "creation" of the object in the real world. Next in line for bithrights – in case no municipal department is directly involved in the birth of the object – would be the department that last authorizes the creation/installation of the structures, or the department that is in charge of managing the activities.

488 Though there are still some gray areas especially around jurisdiction over sewers, which is causing friction between Insula and the Venice department of Public Works.

489 Hart *et al.*, 2004.

490 Another promising example is being developed in Venice (Novello and Sartori, 2004). See also footnote number 498.

491 Craglia *et al.*, 2004 (p. 61) specifically mention a concept very similar to the one presented here. See also the "temple" vs. "triangle" structure discussed in Reeve and Petch, 1999, pp. 155-156.

[change what and how?]

To abide by this tenet, towns would simply need to internalize the concept of informational jurisdictions so that each department can assess its own domain of control and plan how to bring its own turf under a sustainable City Knowledge regimen. The nuance here is that the "parent" office or department is not necessarily the one that ends up carrying the burden of the automating/informating tasks. This is simply the place where the "naming" and "coding" of an object takes place, coupled perhaps with the assignment of appropriate permanent attributes to the newborn.

It could well be that the actual data entry and GIS mapping will still take place at a central repository – like an MIS or GIS department – for larger towns, or even outside of the municipality – for example in regional planning authorities – in the case of smaller towns[492]. In other words, the parent department – after naming the child – may decide to relinquish its care to a "guardian". Yet the paternity would always reside in the department/office that actually oversees the modification of the world we live in. The original parent department would be ultimately responsible for the information about such modifications, not the "guardian" that oversees the actual computerization of the data.

The main departure from currently popular practices is the attribution of a special importance to the exact place of birth and death of administrative objects, so that appropriate jurisdictions can be drawn up using consensus approaches based on the locus of such administrative events. The politics of such a consensus approach would surely be quite intriguing and potentially detrimental to the success of this aspect of the City Knowledge approach, yet there will be plenty of uncontested jurisdictions that are unequivocally already under the sphere of influence of a specific office or department. Of course, we should start organizing urban data within these clear-cut, *de facto* domains before we get into the more controversial ones.

492 See footnote number 437

Atomic Distributed Knowledge

The way in which we have proposed to assign the responsibility over the information relating to the birth and death of each object in a city to the most appropriate department, and not to some single overarching entity, betrays another fundamental tenet of City Knowledge – namely that knowledge should be acquired and organized in a distributed manner and not though a centralized effort[493]. Top-down efforts like that of the Ordnance Survey in the UK[494] are appropriate insofar as the mapping of fundamental topography is concerned, especially when there is a huge backlog of territory that is still unmapped via GIS and unrecorded in computerized databases. But, in the long run, any form of top-down, centralized control over the continuous updating and upkeeping of the information infrastructure is doomed to fail since the real data that captures fine-grained local change originates from the frontlines of local government and not from ministerial headquarters in the nation's capital.

capturing fine-grain at the frontline

Just like ants go about their daily business guided by simple rules that control their individual tasks at the local level – yet the ant colony as a whole displays a macrobehavior of uncanny intelligence – so front offices in our cities and towns micromanage piecemeal change at its finest grain within the civic sphere and in so doing almost serendipitously produce an aggregate effect that translates into macroscopic change over time. Private actions that affect the public realm are always filtered through local branches of government. Only few "big" projects that affect our physical territory are decided and implemented at the state or national level[495]. The majority of change is managed through our local city halls, which is why I propose that the way to create – and more importantly to perpetuate – City Knowledge is through a concerted effort at the municipal level[496], and more specifically at the level of each department or division within a municipality[497].

micromanage local piecemeal change

iteratively maximize informational returns

[gradually accrue data]

Working within the confines of the department, the informational needs of each of the internal operations should be analyzed and a priority sequence of small incremental projects could be planned to make gradual progress towards the final goal[498]. Simple informational tools will be developed, at a very gradual and sustainable pace, to address specific operational needs, while data are gradually accrued through a variety of low-cost methods that are progressively refined and optimized to yield the maximum informational return using the minimum financial and human resources[499].

[minimize cost]

493 In line with Nedović-Budić, 2000, who support the development and maintenance of local databases (p. 87).

494 See footnote 486 and page 222.

495 Interstate highway systems, water supply projects, and federal buildings come to mind, as do military reservations and coastal developments related to navy yards, as well as the national and state park systems.

496 This should be taken to include townships and counties, for the more sparse parts of the world.

497 Tulloch and Fuld, 2001. After Hart *et al.*, 2004, I have been also advising an undergraduate diploma thesis that entails a thorough analysis of "information flow" in the management of permits for the occupation of public space in Venice (Novello and Sartori, 2004).

498 This is not a new approach at all. See chapters 5 and 6 in Reeve and Petch, 1999.

499 Cost savings are reported by Budić (1994) and Nedović-Budić and Pinto (1999), p. 55.

[maximize human resources]

This decentralized approach to the accrual of city knowledge lends itself to contributions by forces outside of the municipal sector. In fact, the work of volunteers[500], students of all grades (K-20)[501], scholars[502], and professionals[503] could be harnessed and incorporated – with proper validation – into this emergent, distributed information system.

what's new about atomic distributed data

The grain of data a city collects has always been controlled by the balance between costs and benefits. The fine atomization that we were able to employ when collecting urban data at the Venice project center since 1988 would have been prohibitively costly even for wealthy western communities until very recently. The middle-out approach suggests that we devolve data upkeep to the front lines, hence implying a distributed architecture for the ensuing information system. Neither of these concepts are particularly revolutionary, but they were simply not cost-effective at the department level 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 hence feasible. Although there will be coordination and synchronization costs, they are beginning to be offset by the benefits of this approach.

[empirical basis]

In Venice, we have demonstrated that sizeable components of the urban realm can be systematically and exhaustively collected, with patience and with proper information design and knowledge acquisition methods[504]. Our inventories have shown resilience to technological change, as we 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 and as emphasized by our plan-demanding cases.

Although we have been operating in Boston only since 1999, we already have achieved considerable success in promoting meticulous comprehensive inventories of Cambridge's curb-side parking regulations and parking meters[505], as well as of Quincy's public buildings[506] and Boston's parking facilities[507].

[similar to other efforts?]

As mentioned, the closest example to this approach, in terms of the distributed and emergent nature of the system is the *Digital Earth* effort[508]. Although it is focused on earth sciences at the planetary scale, it reflects all

500 As the over 200 Earthwatch volunteers who were instrumental in the rapid completion our Public Art projects discussed partially on page 110.

501 Like the over 500 WPI students who came to Venice and Boston over the years. Or the 1000 middle-school children who measured the hydrodynamics of the entire lagoon simultaneously under my direction (Carrera, 1998).

502 As am I planning to do with the emergent transcription system discussed very briefly on page 111..

503 As we proposed to Cambridge to harness traffic reports from consultants. As another example, I think surveyors should be enticed into some submission requirement (by the county-level Registry of Deeds) to make our cadastral system sustainable.

504 See for instance the bridge, dock, public art and canal catalogs in Part II.

505 Cullen *et al.*, 2002; Flynn *et al.*, 2003.

506 Blizard *et al.*, 2004.

507 Allard *et al.*, 2001.

508 Crockett, 1998 and http://www.digitalearth.gov. For more examples of similar systems, see the list of links at http://www.digitalearth.gov/analogs.html.

of the main tenets of city knowledge, making it a true emergent system, albeit at a different scale than my municipal approach. It has the same flavor in terms of distributed cooperation between independent agencies, but a different – much larger – grain.

The Federal Geographic Data Committee (FGDC) seems to promote more of a centralized clearinghouse concept at this time, though there is an overall distributed approach to the collection and organization of the fundamental framework datasets: geodetic control, orthoimagery, elevation, hydrography, governmental units, and cadastral information. Connected activities such as the Geospatial One-Stop and "The National Map"[509] are producing appreciable results with many local initiatives being spawned every month in local areas. Yet these approaches all hover at a scale and resolution that is lower than the urban fine grain that I propose.

Other initiatives from the bottom-up, such as the various Neighborhood Knowledge initiatives[510] and the National Neighborhood Indicators Partnership (NNIP)[511] also resemble my distributed approach, but focus on socio-economic indicators and not on the physical elements that municipalities also require information about in order to conduct routine maintenance, exercise proper management and produce sensible, well-informed plans.

[different/better?]

It is easier to find initiatives that resemble what I propose by looking at examples at the municipal level, though the documentation about the specifics of each city's implementation are hard to track down, making a comparison with my proposal difficult if not impossible. Cities like Vienna[512] or Philadelphia[513] demonstrate some of the more advanced municipal information systems and strategies, though the implementations seem to still betray a dominance of top-down approaches.

My approach is a hybrid that combines the emergent and gradual approaches of the federal efforts like "Digital Earth" and the National Spatial Data Infrastructure (NSDI), with a more bottom-up ingredient similar to the neighborhood data efforts. My focus is at the municipal level and even more specifically at the level of departmental offices. I think that my approach, now that it is technically and economically feasible because of technological developments, promises to be more sustainable since it counts on the finest grain of urban data to produce the higher level information and the second-order knowledge that many of the other initiatives already focus on.

[change what and how?]

The way forward for an interested municipal department would be to pick a low-hanging fruit and begin the process of creating a municipal framework into which to plug urban data as they are collected from now on. The basis of my distributed approach are the aforementioned informational

---

509 See http://www.geodata.gov (last accessed 9/7/04).

510 See http://www.urbanstrategies.org, http://nkca.ucla.edu/, http://nkla.ucla.edu (last accessed 9/7/04).

511 See http://www.urban.org/nnip/ (last accessed 9/7/04).

512 Wilmersdorf, 2003. See also footnote 34.

513 See http://www.phila.gov/mois/index.html.

jurisdictions, so an ambitious office could get the ball rolling as soon as it identified a suitable first project. The thorny issues of coordination of distributed agents, synchronization, and replication that are standard fare in the MIS and IT fields will eventually need to be resolved through the top-down coordination phase of my City Knowledge approach. Meanwhile, the bottom-up, high-resolution data collection can be started at any time, provided that a systematic and exhaustive approach is followed, in accordance to the lessons listed in parts II and III.

For those cities and towns that are already collecting and mapping urban data, the change would be more in the direction of finer grain, richer attribute sets and exhaustive and systematic data collection. For such cities, the next step would be to work on the more advanced aspects of City Knowledge that are described in the sections that follow.

[perpetual updates]

To truly bring each small project to completion according to City Knowledge principles, the information system will not simply include an inventory of all pre-existing assets up to today, but will also include a mechanism for maintaining such an inventory up to date in perpetuity.

The next chapter explains how we envision these updates could take place semi-automatically whenever a change happens in the real world.

Sustainable Updates

Once cities embrace City Knowledge principles and systematically collect and organize data about the various elements that make up our urban world, the next hurdle is to devise methods to keep the information up to date from that moment on[514]. The provision of perpetual mechanisms for updating the data in an information system is one of the distinguishing features of a true City Knowledge application. Such a system not only provides a user-friendly multimedia GIS interface[515] to aid the municipal end-users in their day-to-day urban maintenance, management or planning activities, but also incorporates the data-updating mechanisms in the system from its inception and not as an afterthought[516]. Some of these mechanisms do not necessarily entail purely technological solutions, but often require a combination of technology, together with appropriate changes in policies and procedures. In fact, akin to what was discussed earlier[517], there are at least five main ways to achieve a sustainable level of informational upkeep without massive financial investments:

1. Intercept administrative transactions (e.g. permits);
2. Force contractual updates (e.g. force contractors to return up-to-date information);
3. Change job descriptions to include "informational returns" (e.g. make information updates officially part of the "job" for civil servants);
4. Exploit free or inexpensive labor such as students, interns and volunteers;
5. Budget and plan for periodic updates, particularly for dynamic activities that change over time;

intercepting administrative transactions

As I pointed out at several junctures in this treatise, there are few – if any – modifications to the physical realm that we live in, that are not in some way preceded or accompanied by an administrative act that results in some sort of entry in the public record. The obvious exceptions are private modifications to one's property, when they do not require authorization[518]. The existence of such a paper trail is almost guaranteed to exist in the realm of "structures", whereas conversely it is highly unlikely to exist in the arena of "activities". Unfortunately, the availability of these transaction logs is not routinely exploited as an information source through which a City Knowledge system can be maintained up-to-date[519]. Not only are "births"

514 See Nedović-Budić and Pinto, 1999, p. 58, under "Responsibility".

515 Like the ones shown in Part II and Part III of this paper.

516 In Venice, our *SmartInsula*, *EasyBridge* and *EasyDocks* systems all included updating mechanisms. In fact, the proof of the success of our update mechanisms is that the data these systems contain today are different from what we originally delivered to Insula S.p.A. and the City of Venice in the late 90's.

517 The discussion on page 166 was focused exclusively on the "Catching up with the backlog" aspect of City Knowledge accrual, whereas here we are discussing the subsequent upkeep of the accumulated knowledge, hence the different slant of the section.

518 In Spencer, Massachusetts (and probably elsewhere) these projects are commonly called "ANR", meaning *Authorization Not Required*. Despite the name, many projects that claim to be ANR are still reviewed by planning boards and/or zoning boards in order to ascertain whether authorization is indeed required or not.

519 Despite efforts by researchers such as Coulton *et al.*, 1997.

and "deaths" recorded in these archives, but so are also any subsequent piecemeal modifications, corrections, adjustments and even some of the maintenance performed on the structure – particularly if it is government-owned[520]. Therefore, it seems obvious to me that intercepting existing administrative data streams[521] ought to be the primary means to extract updated information out of the documentation that is associated with acts that require government oversight and are therefore already a matter of public record stored in some municipal recordkeeping system.

Beyond tapping into these administrative records, the operating principle would otherwise be to try to shift the burden of the maintenance of information to outsiders who have an interest or an obligation in the upkeep of such knowledge. The next four sub-sections describe some of these ways.

maintenance-based updates

[contractual obligations and bids]

As mentioned before, there can be many creative ways to incorporate information updates into routine maintenance activities. The tree[522] and light bulb[523] examples are exemplary of these "maintenance-based updates" that should become standard fare in future contractual negotiations and in the language of outsourcing bids.

information-conscious job descriptions

Next the focus ought to be in modifying the Standard Operating Procedures (SOP) internally to include information maintenance. This way internal staff will be made conscious of the importance of information in all aspects of municipal functioning.

free or inexpensive labor

Lastly we ought to look at untapped *pro bono* resources as a final source of low- or no-cost sustainable updates by considering ways to harness the power of volunteers, students, scholars and professionals.

budgeted updates

In addition to these no-cost or low-cost mechanisms, one can also envision dedicating some municipal funds to support additional ways to keep city knowledge current, through focused programs that would fill-in wherever there might be informational gaps left, and validate the data, after these other inexpensive venues have been fully exploited.

Regardless of the method employed – and it may well be a hybrid of the ones above – once these principles are adopted, the imperative will be to never waste any opportunity for updating our city knowledge from that moment forward.

what's new about sustainable updates

As was the case for all of the previous foundation elements, this too is not a particularly novel principle. So-called "lifetime" models of information systems implementation have long included provisions for the updating of the underlying datasets[524].

[empirical basis]

Our own information systems for the maintenance and management of docks (*EasyDock*) and bridges (*EasyBridge*)[525] incorporated screens for the recording of conditions and for the logging of maintenance activities. In

---

520 Some of this record-keeping will be mandatory and may even become standardized by the GASB-34 accounting requirements. See footnote 482 and http://www.gasb.org.

521 Coulton *et al.*, 1997.

522 See page 127.

523 See page 185.

524 Reeve and Petch, 1999.

525 See Part II.

Cambridge, the parking meter collection crew now directly maintains its electronic log book of jams and it – not the Public Works department who installs the meter stand post – is in charge of updating the meter information whenever a head is installed or removed.

[similar to other efforts?]

Of course, everyone expects that data will somehow be kept up-to-date, so traditional MIS literature always includes considerations about information upkeep[526]. The interception of transactions as a means to achieve reliable updates is also not new[527]. Traditional "waterfall" models of information system development[528] always envision a data maintenance and review step at the end of the waterfall.

The neighborhood indicator programs (such as the NNIP and NKLA efforts)[529] have demonstrated that it is possible to tap into statistical or scientific data sources reliably and repeatedly. Stubborn translation issues could be taken care of with the middleware that Ferreira envisioned[530]. Local frameworks[531] could also facilitate the upkeep of the data as the grain gets finer and finer. The GASB-34 accounting mandate[532] and the spreading of asset management tools may soon make the upkeep of data about physical elements of the city more commonplace too.

[different/better?]

The sustainable updates I suggest here differ from other approaches to data upkeep primarily in focus. One view may be that the "low hanging fruit" here are the slowly changing elements of the physical environment so that a simpler system to tap into government data sources[533] may be devised and implemented. Another view may assert that the best return on investment (ROI) is more likely to come from an application where updates are frequent and the process is important to the city. In my view, data maintenance needs to be a fanatical pursuit. There is no point in developing a comprehensive municipal information system if the data are going to be obsolete the moment the system is unveiled.

Another fine distinction is my insistence in leveraging outside self-interest to keep the records up-to-date. Towns already force developers to pay for the services of a planner who will support the town in its deliberations on the developer's project. I therefore suggest that the self-interest of developers could be exploited also to delegate data entry and GIS mapping to them since they are the ones who indeed will be changing the real world out there anyhow.

One potential benefit would be the ability to collect backlog data as well as new updates using essentially the same method. This method may be a composite triangulation of the five paths discussed earlier in this section[534].

---

526 See for example Laudon and Laudon, 1996.
527 See Ferreira, 1998 and 2002.
528 Reeve and Petch, 1999, chapter 3.
529 See the Introduction (pages 11 and 16) as well as later, on page 189 and elsewhere in footnotes.
530 Ferreira, 1998.
531 Tulloch and Fuld, 2001.
532 See also footnotes 482 and 520.
533 See Coulton *et al.*, 1997.
534 Page 210.

[change what and how?]

The course to take in order to establish a solid and sustainable system for reliable data updates needs to proceed opportunistically[535] starting from the more cost-effective updates, as evidenced by an internal assessment of information flows and information sources in standard administrative processes that entail spatial decision-making or analysis[536]. Since we believe that sustainable data maintenance mechanisms are necessary conditions for the longevity of an urban information system, a true City Knowledge system should never be conceived without making provisions for keeping the information current. By including update considerations in the department's data collection strategy it may be possible to exploit possible synergies with the collection of the backlog information so that old and new data can be collected using the same seamless procedure.

Using the five tenets described earlier, the department could revise its requirements[537], modify its forms[538] and generally shift the burden of data upkeep to interested third-parties who may not mind the extra burden as part of doing business with the town.

535 Reeve and Petch, 1999, p. 156; Barr, 1991.

536 As was done in Boston by Hart *et al.*, 2004 and repeated in Venice by Novello and Sartori, 2004.

537 As we suggested in to the Traffic dept. in Cambridge (Gage *et al.*, 2003) and to the Boston Air Pollution Control Commission (Allard *et al.,* 2001).

538 As proposed to the Boston Fire Department (O'Donnell *et al.*, 2002) and to the Boston Environment Department (Hart *et al.,* 2004).

INFORMATION SHARING

Until some degree of sharing is initiated within and among municipal departments, we will not be able to exact the powerful, value-added benefits of City Knowledge that allowed us to quickly and easily conduct many second-order analyses yielding unforeseen plan-demanding results as described in parts II and III. Nevertheless, sharing of the "meatier" datasets, beyond the mere publicizing of the "official" reference codes and the availability of the GIS layers with the objects' positions, ought not be mandatory but rather voluntary. Nobody will be forced to share data with anybody else, unless there is either an institutional mandate or a desire to do so on the part of the rightful "owner".

voluntary sharing

institutional sharing

mandatory sharing of public records

Some forms of sharing will be more or less compulsory, based on pre-existing institutional requirements. For instance, public records access laws[539] and "right to know" would constitute a mandate to share[540]. The most immediate type of institutional sharing is the one that takes place within a department or division. Intra-departmental sharing is a patently obvious form of sharing that ought to take place within municipal organizations, for very apparent reasons. Yet, in reality, the level of sharing that occurs even within small organizations is surprisingly low, despite the intuitive expectation to the contrary. I have briefly mentioned the benefits that the city of Cambridge's Traffic, Parking and Transportation department has reaped from our recommendation that the parking control officers share the daily log of parking meter jams with their colleagues (in the same department) of the meter collection crew[541]. Yet this type of intra-departmental sharing was not happening before our project.

[intra-departmental sharing]

[inter-departmental sharing]

The next level of potentially mandatory sharing could take place between two departments that must communicate information to each other as part of their institutional duties. For instance, the building inspector must consult with the conservation commission about possible restrictions on developments that are near wetlands[542]. In this case, the conservation commission is obliged to share its information with the building department[543]. Likewise, I already discussed how local landmarks commissions must notify the inspectional services (or code enforcement) department whenever a particular building becomes listed as a registered historic property, since different rules may apply in relation to building permits or codes[544].

Intra- and inter-departmental sharing represent the "low hanging fruits" wherein information exchanges could be quickly mainstreamed through a GIS-based, permission-enabled, distributed urban information system.

---

539 Like Massachusetts 950 CMR 32.00.

540 This may also be a hook that state or regional agencies can use to induce smaller agencies at the municipal level to share information with them, so that they in turn can make the appropriate information available to the public at large in accordance with the law.

541 Flynn *et al.*, 2003. See footnote 361.

542 See Hart *et al.,* 2004.

543 Also known as: "building inspector", "construction dept.", "inspectional services", "code enforcement" and others.

544 Hart *et al.,* 2004.

extra-mural sharing

The next step beyond intramural sharing, would entail sharing with “outside” institutions, in situations where contractual mandates exist, or where the mutual benefits make sharing a desirable process. For instance, the electrical utility company (whether or not it is owned by the municipality) has an interest in keeping track of the interference of tree branches with its power lines. Power companies will trim branches of municipal (and private) trees to protect their grids. The Parks Dept. may be in a position to alert the utility company when it detects, through its routine maintenance activities, dangerous situations that may negatively impact the electrical wires. An “asynchronous alert” may take the form of an updated tree record wherein the distance of the tree branches from the closest power line is simply modified as a result of a visual inspection by Parks department staff. This modified record could be shared – through the granting of appropriate field-level permissions – with the electric company which in turn would have its own (asynchronous) built-in middleware mechanism[545] for flagging the “dangerous” distances and thus would immediately recognize the problem and send out a bucket-truck to trim those branches at the first opportunity.

[sharing with public-private companies]

[asynchronous alerts]

[two-way sharing]

The trimming of the perilous branches should then be followed by an update, this time by the utility company itself, of the distance-of-branch-closest-to-power line parameter in the shared municipal tree database, to reflect the change that just occurred through the pruning intervention. This second-round of sharing would thus bring the process full-circle, demonstrating how two-way sharing could be used to leverage these mutually beneficial barters of *quid-pro-quos*[546]. Similarly, within the municipal confines, the Public Works department could also be allowed access to the municipal tree databases to keep track of the damage to sidewalks created by the roots of the same trees.

[informational *quid pro quos*]

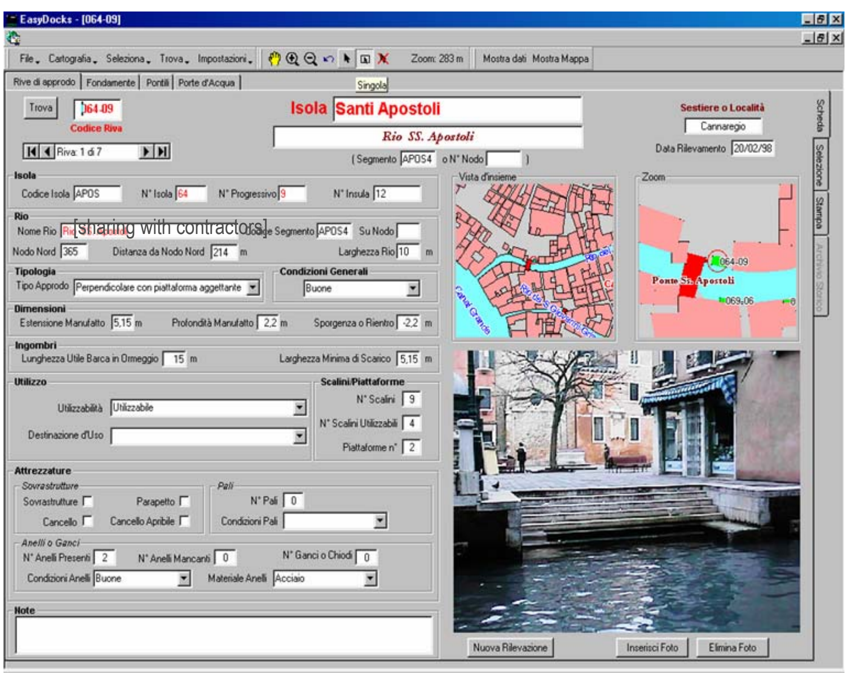

[sharing with contractors]

In Venice, our *Easy Docks* information system (below) for the management of boat docks was immediately shared between the Public Services department and a public-private company (ARTI) that was in charge of the physical maintenance of these municipal assets. The system is still in use at the time of this writing (2004). ARTI (a “contractor”) was expected to update the dock’s condition log as soon as a repair was made.

545 *À la* Ferreira, 1998.

546 See Nedović-Budić and Pinto, 1999, p. 58, under “Incentives”.

Similarly, other outside contractors who are hired to prune trees could be required (by contract) to keep track of the maintenance done to each tree. Contractors could also be obliged to provide updates on the growth and health of the tree. All of these forms of contractual sharing could be achieved through appropriate setting of permissions in shared online databases[547].

[sharing with the public]

States have so-called Public Records Laws that mandate access to public records. The parents of all state laws are the federal Freedom of Information Act and the Privacy Act of 1974. Frequently though, instead of making raw data available to the public, cities prefer to provide pre-screened information – once-removed from the fine-grained original datasets. Thus, a filter is created that allows complete control on the interpretation of the data by precluding independent analysis of the raw facts. Whereas, in the majority of cases, converting data to information represents a major leap in sophistication – one that allows us to conduct second-order analyses that would otherwise be rarely performed due to the difficulty that is generally encountered when just trying to get the basic data together – nonetheless some analyses may be precluded if undigested data are not made available in addition to pre-digested information.

from data to information

For instance, the Environment Department may become interested in determining the energy savings that the urban forest is providing to the city[548]. An adequate energy audit could be easily achieved thanks to the accumulated knowledge about the size of each tree's canopy that would be plan-ready once all of these operations are coordinated around the shared fundamental information references represented by the tree's location (on a GIS layer) and its ID[549]. If the disaggregated, fine-grained tree data were not made available, such calculations would not be possible.

[plan-ready information]

[horizontal and vertical sharing]

The ultimate extensions of this hierarchical sharing scheme involve both horizontal sharing with other cities and towns – probably in the context of regional planning efforts – as well as forms of vertical sharing, both internally in the organization – from the front lines up the management ladder and up the chain of command to the executive branch – as well as among government agencies at different scales, i.e. metropolitan, regional, county, state and federal levels.

[the technology of sharing]

Technically, the act of sharing information among different providers and users can be supported by a variety of client-server or even peer-to-peer architectures. The possibilities run the gamut from file-servers that act as data warehouses and allow sharing through files and network applications[550], to web-GIS applications that allow interactions with shared layers through regular browser interfaces (usually supported by client-side Java applets)[551], or through web-enabled multimedia client applications that

547 As was done for the "sudden oak death" project described in Kelly and Tuxen, 2003.

548 By allowing us to cut our AC use in the summer, thanks to their shading, trees save us money, as they do in the winter by lowering our heating bills thanks to their wind-screening ability.

549 Refer back to pages 125 and ff.

550 Like in the Citrix system in the City of Worcester.

551 See for instance our own demo at www.intelligencesoftware.it/unesco/venezia. See also Kelly and Tuxen, 2003.

can tap into shared layers that are accessible through some internet service and make these updated layers available within a custom, client-side application.

what's new about our information sharing

Not surprisingly, this final aspect of City Knowledge is also well discussed in the literature of MIS, IT, and GIS. Full-fledged geospatial sharing mechanisms are far from becoming commonplace though. Web-GIS prototypes are more and more common, although their effectiveness in day-to-day municipal operations is not so obvious. More commonly, enterprising departments are sharing GIS and Database files through a common file-sharing system, sometimes even through network applications that allow simultaneous access and modification rights to layers and data in real time, without local copies[552].

[empirical basis]

We have dabbled only a little with institutionalized information sharing, since we have scant direct control over it, since our role as academics (and even as professional consultants) is generally that of outsiders. We come into an issue laterally and with no say in policy matters, such as institutionalized sharing. Nonetheless, we have successfully completed some prototyping of potential sharing scenarios both in Venice as well as in Boston, Cambridge and Worcester. We have also recommended specific sharing arrangements that in some cases have been implemented, at least on paper.

[similar to other efforts?]

Interoperability is a common buzzword these days. Coordination mechanisms, such as the ones described in the next section go hand in hand with the challenges of sharing municipal data across jurisdictional boundaries. Structure, process and policies regarding data, responsibility, ownership, contributions and incentives have been suggested as a conceptual framework for making progress on this front.

[different/better?]

My approach is not too different from the latest thinking in this arena, but the main focus, one again, is on leveraging the self-serving instincts of a department and favoring sharing where the quid pro quos are evident and easily achieved for instant gratification. Slowly these sort of success stories – I argue – can make more sophisticated forms of sharing less threatening to the "turf-conscious" individuals who still occupy many municipal positions.

[change what and how?]

Very similarly to what was said in the previous section, this final aspect of City Knowledge should also be left to its own evolution, only with a little bit more awareness of the benefits, possibly advertised or highlighted by self-interested individuals – such as planners and decision-makers – who need the more complicated, articulated, intermixed type of city knowledge that can only be obtained once sharing becomes common-place. Once these second-order advantages are appreciated by the higher echelons of the municipal organization, communicative action theory suggests that these fairly powerful executives and managers can become part of the lobbying group that will promote – more effectively than planners alone – the paradigm shift toward full *information awareness* by treatin City Knowledge as a true municipal infrastructure effective immediately.

552 As can be done with the Citrix systems in Worcester and Cambridge.

The next section explores this latest assumption and other issues related to information sharing.

INTERAGENCY COORDINATION

In concert with the necessary assignment of object-level jurisdictions, a City Knowledge system requires the establishment of a modicum of interagency coordination to enable the sharing of information if and when such sharing should become desirable or necessary[553]. Standards and reference rules will be minimal and as transparent as they are in the similar infrastructure of the World Wide Web, which is the model of "emergent" behavior that this City Knowledge infrastructure hopes to emulate. These municipal standards will also guarantee the reusability of the data and will allow multiple uses by different departments, which may regulate one or more of the following aspects through memoranda of understanding, guidelines or internal bylaws:

[memoranda of understanding]

[guidelines and by-laws]

- ❖ Reference Codes
- ❖ Formats
- ❖ Methods
- ❖ Databases
- ❖ GIS layers
- ❖ Metadata
- ❖ Software applications
- ❖ Network Access

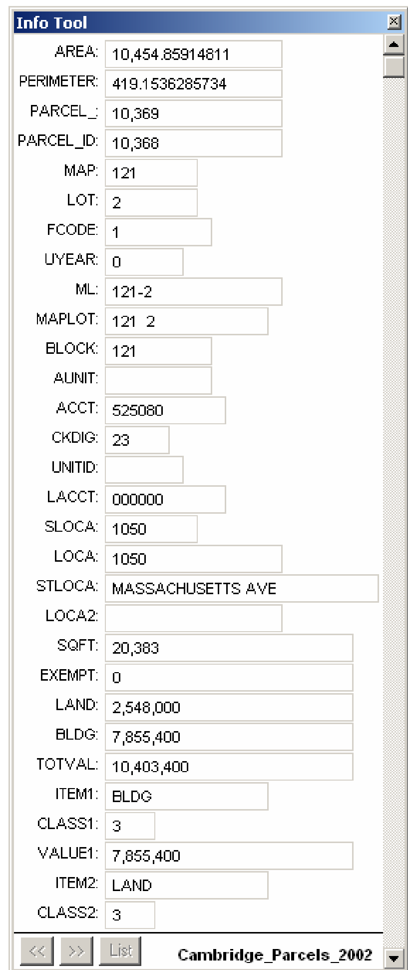

reference codes

Ample space has already been dedicated to the primary form of sharing which revolves around the reference codes, i.e. placetags, that uniquely label real-world objects subject to municipal maintenance, management or planning[554]. Beyond the assignment of jurisdictions to dictate who should be the "parent" to the newborn "city object" and hence have the right to "name the child"[555], there is an overarching citywide necessity to agree on what these identifiers ought to represent (if anything).

There need to be

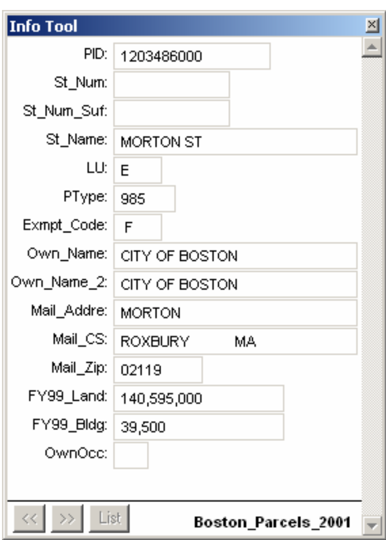

553 See Nedović-Budić and Pinto, 1999, p. 58, under "Data". According to them "interorganizational systems require standards on data models; data formats; data quality; categories of spatial data; contents of specific data layers; metadata; data dictionaries; output requirements; and data transfer."

554 See for instance page 52 and following.

555 See the "Informational Jurisdictions" section on page 201.

some basic agreements about the criteria that the various departments ought to adopt in the process of creating standard codes. In general, all codes should be at least unique, consistent and coherent[556]. For instance, there may be an overall agreement that all codes will be always alphanumeric and the code data type will always be character (or text, or ASCII) and not integer or numeric. This uniformity will facilitate linking, when data types need to be specified or manipulated, so no conversions will be necessary on the fly. Another overall agreement may be that no code should exceed *n* characters in length (probably 20 would be an appropriate number). This simple rule would allow everyone to simply arrange for codes of 20 characters to be set aside as linkable entities in the various department databases. Yet another useful retroactive measure to take is to give each data field a proper, clear name that conveys exactly what that field contains. Indeed, it would be useful to get into the habit to do the same whenever files are named as well.

More importantly, the syntax and semantics of the fundamental codes that will be used for sharing across databases and departments ought to be agreed upon in some sort of conference committee with representatives from a variety of departments.

[legacy codes]

Whether or not new codes are introduced to do away with awkward anachronisms, it is always wise to retain all of the possible legacy reference identifiers that refer to each object, in order to maintain backward compatibility with any dataset that may reference the old codes. Perfect backward compatibility may be impossible due to spatial mismatches between the old codes and the new/updated spatial objects, nevertheless this effort should still be made to allow longitudinal analyses with archival records that, despite their antiquity, may still hold significant informational value for the establishment of long-term trends or for before-and-after comparisons with today's data. Conversely, it is equally wise to stop actively using codes whose meaning or origin is lost to current institutional memory.

[code adoption]

[parent department]

When various departments create databases that link to physical objects outside of their birthing jurisdictions, it is essential that they adopt the precise codes that the "parent" department has assigned to those objects. In this way, sharing will always be possible, regardless of whether it is currently desirable or not. Of course, this transition would need to be coordinated as do many other aspects discussed in this section.

[ancillary multimedia documentation]

[codes as file names]

In addition to linked layers and databases, advanced city knowledge systems will frequently contain ancillary multimedia information, such as photographs, videos, graphs, and sounds that provide additional information or documentation about objects, but are generally not incorporated into either the GIS or the database for reasons of efficiency. The way I prefer to deal with these items that are linked with the GIS and DB through integrated multimedia interfaces[557], is to use, as the file names of the ancillary documentation, the same exact codes that uniquely reference each database

556 These are standard database principles that can be found in any RDBMS textbook.

557 Such as those shown on p. 81 and 101.

record, using suffixes to further differentiate between multiple media attached to the same object/record[558].

[syntax]

[algorithmic code generation]

The syntax of the various codes used as unique identifiers and/or filenames may be the object of a standardization effort within the municipality at a later date. For instance, the city may decide that characters should always be preferred to numbers (even if they are both treated as characters). A possible standard may recommend creating codes from the concatenation of a variety of fields that are already in the database to identify the object. As an example, we used the combination of island code and a sequential number to label each sewer outlet in Venice[559]. The specific syntax used within each individual application need not be agreed upon by the entire city, but generic syntactical "rules" may be part of a citywide standard nonetheless.

Furthermore, smaller committees of interested parties ought to get together to share the specific syntax of items of common interest. Even though the "parent" department has full control over the naming, it may be wise to coordinate the coding syntax with departments that are likely to interact with a specific category of objects. Outside of these object-specific syntax committees, the citywide standardization process would only address general "ways" or "criteria" to adopt in the definition of the syntax of a code. Maximum freedom will remain in the determination of the exact syntax for each specific category of objects within each individual jurisdiction.

Sophisticated algorithmic procedures may be devised to generate codes automatically from some implicit or existing parameter that already contains the seed of uniqueness that is necessary to establish a successful referential framework. One such scheme would entail exploiting the singularity of centroid coordinates (for non-overlapping objects) to construct a composite, interlaced unique code that mixes the X and Y coordinates (regardless of the projection used)[560].

[558] The two pictures that accompany each bridge record are labeled with the bridge code (e.g. GOLD) + a suffix of A or R to represent the "Arch" or the "Ramp" photo (GOLD-A.jpg and GOLD-R.jpg). We have used a similar system for wellhead pictures, where each well was photographed on average 10 times, from different angles. Of course, if dozens of pictures were needed for a single object, one may eventually revert to numeric sequential or time-stamped suffixes.

[559] As described on page 90

[560] The reason for the interweaving, as opposed to the simple concatenation of the two coordinates is so that an alphabetical ordering of the objects according to the interwoven ID would not privilege one coordinate over the other, but would combine the two and organize the objects more or less according to their proximity across both dimensions. The disadvantage is that code generation is much more complicated, requiring a "small" program versus a single concatenation operation.

[automatic codes]

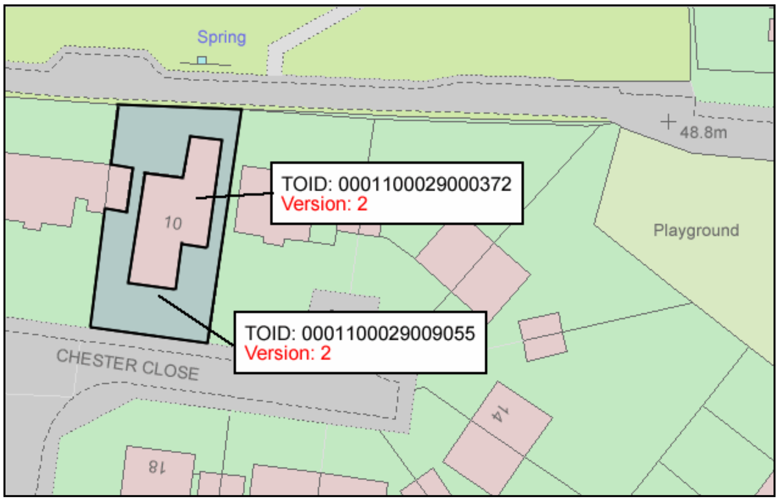

The Ordnance Survey in the United Kingdom has begun to assign automatic codes to all objects that it maps in Britain, by attaching a 16-digit TOpographic IDentifier (TOID) to each mapped object. This rather "dumb" assignment of codes in an algorithmic manner has the advantage of creating an automatic reference framework through which all other agencies can refer to the "real" elements that they are each responsible for. Although some objects may actually be composed of several TOIDs, this approach has the implicit advantage of facilitating sharing.

[code semantics]

["talking" vs. "dumb" codes]

Some of the syntactical rules adopted by a city may embody semantics of one sort or another. For instance, a suffix appended to a code may indicate a sub-partition of the object. A prefix, on the other hand, may serve as a way to identify the "group" of objects that an item belongs to. A syntactical standard accompanying such semantic breakdowns may impose the use of hyphens (instead of underscores, for instance) to separate prefixes and suffixes from the body of a code.

[numeric IDs]

Traditional RDBMS literature[561] recommends the use of numeric IDs that are much more computer-friendly since they lend themselves to effortless sequencing and other algorithmic computations like indexing and sorting. The ease of spatial operations makes "talking", spatially-explicit codes just as easy to maintain with algorithmic precision. At any rate, a compromise can probably be arrived at by providing both a computationally-efficient number and a mnemonically effective code for each object.

Finally, the suffixes that we frequently use to break down sub-units of a bigger whole are always chosen with some logic that embodies semantics. The same is true of the suffixes of file names that refer to images and other media associated with our data. The file name 101_L.jpg refers to the photo of the Lid of wellhead number 101.

At any rate, leaving aside my personal preferences, all I am proposing here is that these issues be discussed and agreed upon across the entire municipality in order to create a standardized framework for future citywide sharing.

formats

In addition to semantically mnemonic and syntactically consistent codes, the formats of the datasets, of the GIS layers, of related files and of any ancillary data ought to also be agreed upon at the city level. Agreement on formats may at first be limited to acceptable and unacceptable file types that departments should uniformly adopt and reject, respectively.

---

561 Such as Ullmann and Widom, 2001.

methods

Departments may also decide to standardize the *methods* that field crews or professional consultants will adopt for the collection, archival and presentation of data, information and knowledge.

databases

A little deeper level of standardization may relate to the structure of some of these files, dictating for instance that all Access database files contain tables with names preceded by a numerical sequence number, to organize the contents of the MDB file. Database fields may also be the object of some interagency agreement. Some databases may become standards as a whole and be incorporated into high-level framework datasets.

GIS layers

Some core sets of GIS layers may be standardized, along the lines of what the FGDC is doing for "framework" data[562] so that there may be even compatibility across town boundaries – with abutting towns – even across state lines or vertically, from one level of government to the next higher (city to state to federal).

metadata

[extramural and vertical sharing]

The exact structure of the metadata used within a municipality is one of the aspects of City Knowledge that ought to be standardized on a citywide (or even state or federal) basis. The management of evolving versions of both layers and datasets could be achieved through a strict abidance to the metadata standards that are slowly emerging in the GIS industry. As of this writing, it appears that the most useful approach to this issue would be an interdepartmental agreement on what "subset" of the all-encompassing FGDC metadata standard to adopt[563]. At the very least, in the beginning, the system may rely on the simplest file-system metadata that operating systems already provide in the form of creation/modification dates, owner, permissions, and file size. Basing the municipal metadata on an existing standard will make possible the next level of sharing, beyond municipal walls, with other cities and towns, or beyond the municipal level of governance, with state and federal agencies.

software applications

It may be advantageous for some towns to force the use of specific software applications (e.g. use Mapinfo instead of ArcGIS, Oracle instead of Access) to facilitate sharing and more importantly to cut the cost of software support. Although this approach would certainly save money in the long run, it may engender resentment on the part of those who are forced to switch to an unknown package and may also preclude some "innovation" from happening. Such a level of standardization is no longer necessary thanks to today's highly interoperable software packages[564], so this standardization may not be advantageous except in terms of software licensing and support costs.

network access

Assuming that the data will at some point be shared through a network, there needs to be some coordination about network access, regarding passwords, permissions, quotas etc. These agreements could be coordinated with the department that directly manages network operations so that everyone can access the system remotely, with the appropriate level of read/write access.

---

562 Tulloch and Fuld, 2001. See also the FDGC web site at http://www.fdgc.gov.

563 As MassGIS is trying to do in Massachusetts (see http://www.mass.gov/mgis).

564 Mapinfo can read and modify native ESRI shape files, for example. Similarly, SQLServer is capable of reading Access files.

what's new about interagency coordination

Standardization is a necessity that emerges naturally when multiple actors are trying to cooperate toward a common goal. My approach is hardly new or different from the myriad of examples in typical municipal MIS, GIS and RDBMS practices.

[empirical basis]

Both in Venice and in Boston, we developed and successfully deployed infrastructural reference codes to uniquely identify key physical elements of the built environment, such as canals[565], bridges[566], docks[567], pieces of public art[568], parking facilities[569], underground storage tanks[570] and so on. On both sides of the Atlantic, we have also structured spatial frameworks for the archival of dynamic data about such activities as traffic[571], demographics[572] and economic vitality[573]. Many of our standard codes have become *de facto* standards in Venice[574] and in the U.S.[575].

We also standardized photo formats (JPG) and methods (landscape vs. portrait). We structured our Access tables internal to the MDB files in predictable ways and we even developed a set of fundamental GIS layers, similar to the local framework layers proposed elsewhere. We have not, alas, dug deep into the issue of metadata beyond the simplest file system metadata and little more. We have settled on standard software applications, specifically Mapinfo and Microsoft Access for now. We have also shared standard layers internally through passwords and a "secret" web site on WPI's server. A web-GIS prototype is being updated to improve on the promise of the interactive, password-protected system that is already available on the internet.

565 See part II, especially the chapter entitled "The Venice Inner Canals Project", starting from page 44.
566 *Idem.*, see also section starting on page 99.
567 *Idem.*
568 See page 112.
569 Allard *et al.*, 2001.
570 O'Donnell *et al.*, 2002.
571 See for instance Carrera, 1996, 1997, 1999a; Gage *et al.*, 2003; and Farmer *et al.*, 2004 to name just a few. See also page 70 and following.
572 Hamir *et al.*, 2004.
573 Jajosky *et al.*, 2004.
574 The canal, bridge and dock codes that everyone uses in Venice are essentially the ones we developed (Carrera, 1999d).
575 The curb-side regulations in Cambridge, for example, are now identified by our codes (Cullen *et al.*, 2002 and Flynn *et al.*, 2003).

[similar to other efforts?]

Our standardization efforts resemble the initiatives that are being conducted at the national level by the FGDC and Digital Earth and at the state level by the like of MassGIS in Massachusetts. Technically, the metadata standards of the FGDC and the interoperability progress being made by the Open GIS Consortium will probably create the premises for fruitful advancements in these areas[576]. At the local level, there are some movements to define standards for indicators[577] at the NNIP and NKLA. Some cities are more advanced than others at defining common practices and standards and are more effective at coordinating the mapping and data collection efforts of a variety of departments, though the approaches are often either centralistic or strongly hierarchical and implemented from the top down.

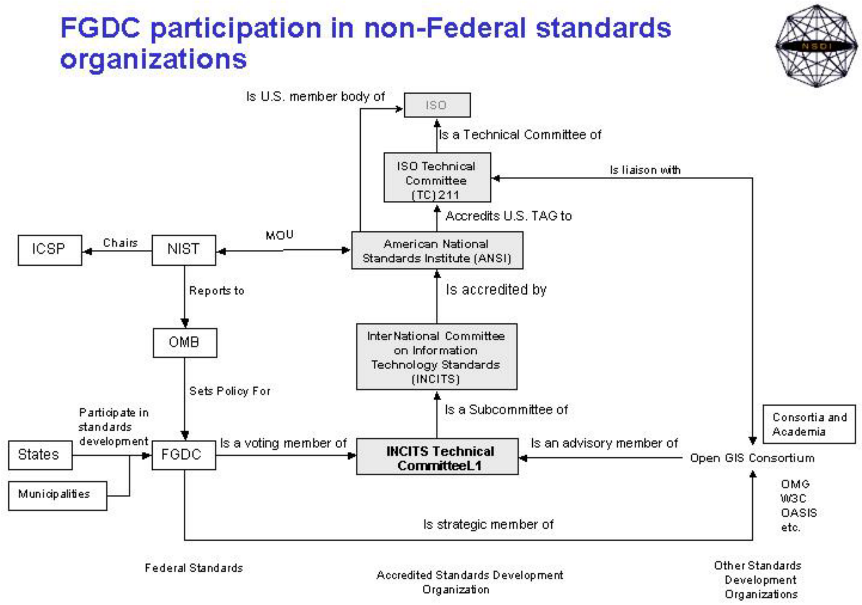

Despite all of these efforts, the systematic codification of all homogeneous GIS objects at least for the fundamental framework layers is still not standard practice even in technically savvy and fairly wealthy cities of the developed world, such as the ones discussed herein. Although the concept of unique reference codes is common knowledge and even common practice in the MIS and RDBMS fields, many GIS layers out there are still CAD-like spaghetti files that look good but are not useful for spatial data archival and analysis. Even when selectable objects are mapped in GIS (instead of the lines and points of spaghetti-maps) they often lack the all-important key ID[578]. The same goes for all of the other aspects that are amenable to standardization, such as formats, methods, databases, GIS layers, metadata, software applications and network access for data sharing.

[different/better?]

My proposal is not fundamentally different from these existing examples, but once again it focuses at a local spatial scale, where high-resolution urban data can be collected from the middle-out. Even though overarching standards are an eventual necessity in a full-fledged City Knowledge system, they do not have to be forced upon reluctant municipal offices from the top. The need to coordinate should be allowed to rise spontaneously from the natural, organic evolution of the municipal information systems and from the inevitable need to eventually interact across divisional or departmental boundaries, with other offices of the same

576 See for instance http://www.fgdc.gov/standards/related_activities.html (last accessed 9/9/04).

577 Sawicki and Flynn, 1996; Coulton *et al.*, 1997.

578 For example, most municipal buildings layers do not have meaningful identifiers that can be linked to other datasets.

or even of a different department, not to mention the possibility of vertical aggregation toward the executive level of municipal affairs or to the state or federal level of government.

Once the usefulness of standards is appreciated first-hand by the font-line offices, standardization will become standard fare itself. Coordination will become second-nature as its advantages become obvious to more and more municipal practitioners.

[change what and how?]

In keeping with the "disjointed incrementalism" of our collective "muddling through"[579] and the opportunistic[580], emergent attitude of my middle-out approach, I do not have a silver bullet to offer to municipalities who may want to engage in standardization practices across departments. Instead, I recommend that, after having embarked in the fine-grained data collection and updating mechanisms described earlier, each department should explore intra-departmental sharing necessities and subsequently inter-departmental interactions that routinely occur in daily administrative processes[581]. Once these links are established and analyzed, I forecast that overarching standards will become blatantly useful and hence will have much more of a chance of achieving "stickiness", which is one of the pre-requisites for reaching a tipping point[582].

The five foundations of City Knowledge: jurisdictions, atomized, distributed data acquisition, sustainable updates, sharing and coordination, together with the middle-out approach, have demonstrated great potential for bringing about a comprehensive City Knowledge system, as the numerous examples in Parts II and III clearly showcased. In this Part IV we have looked at the pros and cons of City Knowledge and laid these foundations.

The following Part V closes this dissertation by proposing a possible sequence of actions to get towns from point A (interest in City Knowledge) to Point B (fully implemented City Knowledge system) and beyond.

579 Lindbloom, 1959.
580 Barr, 1991.
581 As we did with Hart *et al.*, 2004 and Novello and Sartori, 2004.
582 Gladwell, 2000.