Pergamon

Comput., Environ. and Urban Systems, Vol. 21, No. 6, pp. 377–391, 1997

Printed in Great Britain
0198-9715/98 $19.00 + 0.00

PII: S0198-9715(98)00009-X

TOWARDS A GEOGRAPHY OF GEOGRAPHIC INFORMATION IN A DIGITAL WORLD

Michael F. Goodchild
National Center for Geographic Information and Analysis, and Department of Geography, University of California, Santa Barbara, CA 93106-4060, U.S.A.

ABSTRACT. *Even though distance is much less of an impediment to interaction in a connected digital world, decisions must still be made about where to store information, and how to search for it. Central facilities location theory provides a framework for discussion of libraries and information stores as services to a dispersed population, based on certain basic behavioral principles. A revised central place theory is presented based on assumptions about behavior in a digital world. It predicts very different patterns of location for information than we observe today, depending on the degree of geographical variation in levels of interest. Massive changes are also under way in the production of geographic information, as production and dissemination shift from centralized to local. Libraries of the future are predicted to emphasize special, locally centered collections. A geographic data set is most likely to be served from locations within its geographic footprint.*

INTRODUCTION

The human population is spread over the surface of the planet in a complex pattern characterized by extreme variations in density. Further complexity is evident when the population is segmented on economic, ethnic, and other demographic bases. Thus providing services to this dispersed population is a nontrivial task, especially when service must be provided from central locations in the interests of economic efficiency.

Central facilities location theory is a well-defined field within operations research, economic geography, and related disciplines (Love, Morris, & Wesolowsky, 1988). It seeks to design optimal locations of services, and allocations of demand to them, under varying assumptions about system objectives, economies of scale, and consumer choice. It is in many ways an outgrowth of the much earlier central place theory (Berry, 1967; Christaller, 1966; Lösch, 1954), which sought to find normative patterns of service provision under much more restrictive assumptions. To cut a long story short, the predictions of central place theory that central facilities will position themselves in certain regular geometric patterns turn out not to be true in reality (Berry, Parr, Epstein, Ghosh, & Smith, 1988; Goodchild, 1972, 1992, p. 149), because the assumptions underlying the theory are much too restrictive even in the most ideal geographic environments. Efforts to free the theory from some of these assumptions, by making use of the power of the digital computer to evaluate alternative patterns in more complex geographic environments, began in the 1960s (Rushton, 1972), and culminated in the location–allocation models of applied operations research.

A large set of services fall within this framework, including such private sector services as retailing (Goodchild, 1984), and such public sector services as emergency facilities, schools, government offices, and recreation facilities. In all of these cases, a service must be provided to a dispersed population from a few central facilities. Models variously capture patterns of consumer choice, availability of sites, existence of prior facilities, and the economics of operation of the central facility. The results of modeling can include locations of facilities, allocations of demand to facilities, the number of facilities, and their sizes.

Libraries certainly fall into the category of central facilities serving a dispersed population. In some cases, libraries serve the general population; research libraries serve special populations. In the latter case, demand is often largely confined to research institutions, where the population requiring the service of a research library is concentrated. Libraries provide access to information, largely in the form of printed text, but also including maps, music, photographs, and other information formats and media. A library's primary function is to place the medium carrying this information in the hands of the user, and this function is supported by ancillary services such as cataloging which allow users to find the information in the library's collection, and circulation which ensures that the medium returns to the library for use by others. Libraries also provide a host of less obvious functions, sometimes serving as community centers, sites for training in various aspects of information retrieval, and as environments conducive to scholarly research.

A system of central facilities evolves under given patterns of consumer behavior, the technology of transportation, economics of service provision, and consumer demand. When any of these change the system attempts to adjust; facilities are added or deleted, or moved, in response to these changes, as the system works towards a new optimum. Several hypotheses regarding this process of adjustment have been advanced and tested (Berry et al., 1988, pp. 103–113). In the case of libraries, the transition to digital information handling is in the process of engendering changes in many aspects of the central facilities model, including access (transition from physical access and delivery of media to access through electronic networks and delivery of bits), economies of scale (physical libraries replaced by digital servers), and consumer behavior (consumers have increasing numbers of choices). The purpose of this paper is to explore the implications of these changes for libraries, within the context of central facilities location theory, and thus to anticipate the geographic restructuring that can be expected to occur; in addition, the paper focuses on the complications that result when the information being handled is itself geographic in nature.

The paper is laid out as follows. The next section provides a brief review of the principles of central facilities location theory; full expositions of the theory can be found in many introductory texts in economic geography. This is followed by a discussion of the concept of the digital library, and its technical feasibility. The subsequent section examines the special case of geographic data, the impact of geographic information technologies on its production and dissemination, and the extent to which geographic information dissemination fits the model of the digital library. Then the concept of information of geographically determined interest is introduced, as a principle of fundamental significance to the theory of digital library location. Finally, the proposed theory is summarized and the likely geographic impacts of technological and economic changes in the transition from physical to electronic dissemination of information are discussed.

CENTRAL FACILITIES LOCATION THEORY

There is a vast literature on the topic of central facilities location, and its classical precursor of central place theory. No attempt is made here to synthesize this literature; instead, only those parts are included that are directly relevant to the problem at hand. Many of the complications discussed in the literature are ignored in the interests of brevity.

Goods and Services

At the core of central facilities location theory is the concept of a good, a physical entity provided by the central facility to the dispersed population. Because central facilities often exist to provide less tangible services, in this discussion the concept of good is assumed to include physical goods, such as books provided by libraries, and also their digital equivalents. A given good is obtainable from any of a set of facilities, and in classical central place theory competition between facilities is assumed to force all to offer exactly the same variant of each good at the same price. In practice, however, goods can be defined at various levels of aggregation, from individual grocery items to the complex of services provided by grocery stores. Thus a key issue in any discussion of libraries as central facilities concerns the definition of good. For the purposes of this discussion, a good will be defined at the lowest level of aggregation: the book or volume in the case of the traditional library, and the *information-bearing object* (IBO) in the case of the digital library. An IBO might be a digitized map, a digitized photograph, an article from an electronic journal, or a digitized book. In each of these cases the library's services are oriented toward storing, abstracting and describing, finding, retrieving, and delivering the IBO as an atomic information entity. It is assumed that the IBO also includes technical information concerning format, handling, and so forth.

The services provided by traditional libraries clearly go far beyond this. In some cases, such as the library's role in functioning as a community center, the impacts of technological transition may be negligible. This paper focuses strictly on the library's role in providing access to IBOs, and on the associated locational implications of technological transition. In a digital world it is clearly possible for the information server to be located anywhere; in the context of access to IBOs, therefore, the term "library" must be interpreted loosely, and does not necessarily imply a building that we would recognize as a traditional library.

Consumer Behavior

Consider a set of central facilities providing some good, indexed by j, $j=1,...,N$. This good is required by a set of consumers, indexed by i, $i=1,...,M$. Spatial interaction models (Fotheringham & O'Kelly, 1989) predict the propensity for consumer i to choose the offering of facility j. Such models include a factor x_{ij} describing the disutility of traveling from i to j to obtain the service; the factor is known as *impedance*. It is often an increasing function of the distance to be traveled, and captures the costs and other effects of having to travel to a central facility. Demand may be *inelastic*, if consumers insist on consuming irrespective of the conditions under which the good is offered, including the distance to the offering; or *elastic*, if consumption is dependent on conditions.

In classical central place theory, x_{ij} is assumed equal to shortest distance. Because competition has forced all providers to offer the identical good at the same price, it is assumed that consumers will always obtain the good from the nearest offering. Consumers will only consume if the closest offering is within a maximum distance known as the *range* of the good.

A key assumption of central place theory is that consumers are fully informed. Shoppers know about the goods being offered at central facilities, their prices, the effects of having to overcome distance, and so forth.

Economies of Scale

The central facility operates under scale economies, which define the cost of providing a certain volume of service, equal to the total service demanded by the consumers who choose to obtain service from this provider. In classical central place theory, the provider is unable to operate until demand exceeds a level known as the *threshold.* Moreover, enough providers will enter the market to ensure that each operates exactly at the threshold and there are no excess profits. Demand is assumed to have uniform density; it follows that facilities will locate at the nodes of a hexagonal grid, the geometric manifestation of the assumptions of the theory.

Agglomeration Economies

In some cases the concept of a good is well defined; for example, in the case of fire stations it is clear what service is provided by each facility, and there is no overlap between this and other services. In retailing, however, a good such as groceries may be provided by several different types of retail outlet, and some of these may provide other services as well. Such overlap of services is of course a major element of retail strategy; but it can make modeling much more difficult. In spatial interaction models, for example, it can be difficult to enumerate all of the possible outlets of a given service, in order to model choice between them.

Consumer behavior drives the agglomeration of services. A range of goods is offered in shopping centers, for example, in response to the consumer's preference for multiple-purpose trips and one-stop shopping. Offerings of the same class of service often agglomerate (the "auto-mall" of many car dealers, the street of many jewellers) for the same reasons.

In central place theory, it is assumed that services are discrete. Furthermore, an ordering of services by threshold exists: *low-order* goods are those with the lowest thresholds, and consequently the greatest number of offerings; *high-order* goods have the highest thresholds, and the smallest number of offerings. In central place theory, offerings of goods agglomerate, such that every location (town) offering a good must offer all goods below it in the ordering. Because each level must form a hexagonal grid, and because there are limited ways in which hexagonal grids of different densities can stack on top of each other, it follows (Berry et al., 1988) that agglomerations of services will occur only at certain discrete levels where the number of offerings at one level is related to the number at the level below it by a factor k constrained to the values 3, 4, or 7.

Central Facilities and the Library

As noted earlier, it is appealing to conceptualize the library as a central facility providing service to a dispersed user population. Like retailing, almost all traditional library services have required the user to travel to the library, with the obvious exceptions of mobile libraries (but see Berry et al., 1988, chapter 5, for an example of the literature on mobile facilities as a response to specific conditions within the framework of central facilities location theory). Two classes of service are common: that of the public library, providing a service demanded by the general public, and that of the research library, serving a much narrower base. In this paper we focus on the second, although many of the conclusions appear to apply also to the first.

As noted earlier, the demand for the services of a research library is highly nonuniform, and most dense in the immediate environs of a research institution. Research libraries are in part a cause and in part an effect of the distribution of researchers. Demand falls off rapidly with distance, although some users may be willing to travel long distances to obtain service. There are clear scale economies to a research library, since the marginal cost of providing increased service is a diminishing function of the volume of service.

The goods provided by a research library are far from uniform, however. The size of the collection is clearly highly variable, ranging up to order 10^7 volumes, while the service provided to undergraduates at a typical campus is defined by a much smaller number of books. Nevertheless, if we regard the individual book as the good, the assumption of central place theory, that services can be ordered and that higher-order services subsume lower-order, still holds to a degree: a large research library almost certainly contains all of the books of a small research library, with the obvious exception of special collections. Various conventions help users determine whether a given volume is available in a given library. A major research library, for example, will be assumed to contain copies of all journals and books; if a user determines that this assumption is false in a given instance, then it is normally possible to make use of an ancillary service to obtain the item from another nearby library. Thus the assumption of central place theory that the user possesses complete information is in effect true; and users are normally able to obtain all services from the closest facility, again in conformity with the assumptions of the theory.

THE DIGITAL LIBRARY

If we assume that the average volume in a research collection contains order 10^6 text characters (order 10^5 words), then the text in a library of 10^7 volumes amounts to about 10 terabytes. Today's 10 terabyte file stores occupy order 10^3 cubic meters or less; thus it is possible for the text in a large research library to be stored in digital form in the space of a large room. The Internet provides access to remote servers to an increasing proportion of the population, either at home or through access points in schools, libraries, and elsewhere. Thus it is technically possible to serve the entire contents of a major research library in digital form to any user of the Internet. This argument of course ignores that portion of the library's contents that are not in the form of text. Although virtually any form of information can now be digitized, some forms, such as imagery and video, require much larger volumes of storage. It also ignores the question of digitizing cost, although almost all *new* acquisitions of text have likely been in digital form at some point in their development. Finally, it ignores the vital issues of inequity of access to the Internet and the protection of intellectual property. Despite these arguments, however, library services are increasingly available in digital form, and that availability is likely to increase in the coming years.

In this paper, we assume that the granularity of information in the digital library will be roughly equal to that in the traditional library: digital servers will serve digital copies of books and journals just as the traditional library has served analog copies. In principle, however, what is actually served by the library may be either a union or subset of IBOs: for example, a user may request only a single chapter of a book, or several volumes of a journal.

Within the context of central facilities location theory, the digital library is a very different entity from the traditional physical library. Most obviously, the effective impedance of distance goes to zero, since in principle access to a given library from a given location is independent of the distance between them. In practice, however, the Internet still shows a distance deterrence, although distance on the Earth's surface is a poor predictor of its value. The *bandwidth* available between a client at i and a server at j at time t is a complex function, as is the *latency* or delay experienced in obtaining a given service. Very little is known about these functions, since they vary through time and are frequently affected by upgrades to the connections of the Internet. Nevertheless it is possible to make certain generalizations. First, both bandwidth and latency depend on the political jurisdictions containing i and j. Anecdotal evidence suggests, for example, that the bandwidth between the U.K. and the U.S. drops sharply in the middle of the U.K. weekday morning and recovers in the early evening. Bandwidth may also be affected by the networks to which i and j are connected: it may be much higher if i and j are connected to the same subnetwork (Hayes, 1997). Finally, bandwidth depends on the local connections of i and j. A large library server will almost certainly have a "fat pipe" to the Internet, such as a dedicated T3 line, at order 10^7 bytes/sec; a library user may be connected with nothing more than a telephone modem, limiting bandwidth to order 10^3 bytes/sec. All of these effects imply that the impedance associated with a connection between i and j in the digital world is still determined to some extent by geographical distance, although much less clearly and strongly than when the user is required to travel to the library.

In the provision of traditional services, the need for duplication of offerings at every outlet is ensured by the impedance of distance (the range of the good in classical central place theory): it is not possible for every U.S. citizen to travel to New York to shop. When impedance goes to zero, or range goes to infinity, it becomes possible for a single outlet to serve all demand, subject of course to the outlet's ability to scale to the required volume. In principle, then, in a perfectly connected Internet, one copy of each IBO is sufficient if mounted on a server that is capable of handling the demand (this is the digital equivalent of the good that is obtainable from only one place on Earth, but is of sufficient value to consumers to have an unlimited range). Because the Internet's connectivity is not perfect, some agencies resort to *mirroring* a site at other locations, notably in other countries, to reduce impedance. Yet while duplication of volumes in every library was a necessity of service in traditional libraries, in the digital library duplication can be the exception rather than the rule. Instead of order 10^3 copies, and a system that measures success by the degree to which each library duplicates the collection of the largest (see, for example, the rankings produced by the Association of Research Libraries), the digital world offers the exact opposite, where success is measured by how few copies are needed. Moreover, service can be provided from anywhere to anywhere.

Although research libraries are typically monolithic organizations, some degree of specialization occurs. Special collections were mentioned earlier. It is also common for segregation to occur between the sciences and the arts; law collections are often maintained separately, as are music, map, and imagery collections. In some cases these separations are driven by physical considerations; maps and music, for example, require special storage. In other cases the services of the library fail to scale, and segmentation reflects the greater efficiencies that result when there is no economy to be gained by integrating circulation, cataloging, or collection-building services over different information types or disciplines. In other cases separation may reflect distinct sources of funding, or desire for control. In classical central place theory, retail services are similarly segmented within the overall agglomeration of a market town or shopping center, and even to a lesser degree within a department store. Nevertheless, the consumer's desire for "one-stop shopping" and price comparison ensures that services will agglomerate spatially, while remaining functionally separate. Also, economic competition ensures that every offering of a good will be essentially the same.

With perfect Internet connectivity, it is possible for library demand to be satisfied by a distributed system of servers, provided every IBO exists somewhere on at least one server. Specialization is possible, with related IBOs being collected on specialized servers. A server might offer all of the writings of its author custodian, or all of the information gathered by a researcher in support of a particular project, or all of the information possessed by a particular agency, or all of the information an agency is required by its mandate to make available. In effect, this is the world of instantaneous travel by the consumer to highly specialized retail stores located anywhere on the surface of the planet.

Such an arrangement would run counter to two assumptions of central place theory, however: first, the consumer's preference for one-stop shopping, and second, the consumer's lack of complete knowledge of the location of every IBO. Both problems are addressed in part by today's Internet search engines. A single IBO mounted on a server is the digital equivalent of a library with only one book, unlikely to be found by anyone not specifically directed to it. But the search engines find and analyze the contents of sites, and build a primitive equivalent of a catalog, based on the occurrence of words and phrases. Thus the agglomeration economies and favorable consumer behavior that result from locating many types of merchandise in one store, or many stores in one shopping center, are echoed in the virtual agglomeration of information provided by a search engine, which provides a "one-stop-shop" in the form of a primitive catalog to IBOs that may be widely dispersed and otherwise unrelated.

Three types of stakeholders drive the existence and scale economies of search engines. From the perspective of the provider of an IBO, the search engine serves to increase access at no cost. It is in the interests of the owner of the search engine to provide as large a service as possible. The owner receives income from advertisers, who benefit in turn because of the increased business that results from being prominently displayed by the search engine. Finally, the user benefits because the search engine provides a primitive catalog to the Internet at no cost to the user.

Despite the apparent placelessness of the Internet, current technology requires information to be served from somewhere and delivered to somewhere. Heisenberg's uncertainty principle notwithstanding, at geographic scales a bit always has an associated location in real geographic space. But hyperlinks and uniform resource locators (URLs) allow an IBO to appear to be in many places at once, since the knowledge of a URL is in effect equivalent to having the information itself, assuming perfect connectivity. To achieve the behavioral desiderata of one-stop shopping, it was necessary that a variety of stores actually be present in a shopping center. In the digital world, the search engine achieves the equivalent by possessing only the URLs of IBOs in its server; perhaps the digital equivalent of visiting a shopping center containing only catalog stores.

GEOGRAPHIC INFORMATION PRODUCTION AND DISSEMINATION

Consider the traditional arrangements for the production and dissemination of topographic maps. The crafting of a topographic map from inputs of air photographs, field interviews, and surveys is a complex and highly skilled task, requiring a high level of capitalization; printing requires specialized presses. Topographic mapping is also of strategic value. Thus the traditional solution has been to take advantage of economies of scale, and to deal with the problems of scarcity of skilled labor, by concentrating production in a few, highly centralized facilities, such as the National Mapping Division of the U.S. Geological Survey (USGS).

The transition to digital information technology has changed the economic and technical conditions underlying these arrangements. The equipment required for map drafting is now potentially available in a large proportion of households and offices. Digital cartography requires far fewer and less sophisticated skills (arguably to the detriment of the product, however). Even photogrammetric processes can now be emulated digitally, using readily available digital imagery as inputs. The digital transition has largely removed the scale economies, high fixed costs, and labor specialization characteristic of traditional production, and reduced the costs of printing and dissemination effectively to zero. Production and dissemination of geographic information can now be organized economically at the scale of the individual farmer, local government, company, or agency.

These changes mirror similar ones affecting the publishing industry generally, or likely to affect it in the near future. Traditionally, editing and printing of books have been skilled tasks with substantial economies of scale, justifying agglomeration of the publishing industry into large firms in central locations. Electronic communication and production have allowed the industry to disaggregate to some extent, and to support increasing degrees of telecommuting. But the advent of the Web has opened the potential for a much more massive restructuring, in which the entire concept of publishing can be rethought. In a world in which the technical potential exists for anyone to write, publish, and disseminate, the role of the publishing industry is likely to be very different.

The institutions responsible for central production and dissemination of geographic information are already beginning to rethink their roles. Recent statements by and on behalf of the USGS, for example (National Research Council, 1993; USGS, 1997), identify a future role of standard-setting, coordination, and research as the actual production and dissemination operations are increasingly disaggregated, particularly to lower levels of government. The future of the national mapping program is seen as a patchwork, held together by national standards, in which detailed information is produced by local governments, and the national government's role in production is limited to coarse coverage of residual areas. The Natural Resources Conservation Service (NRCS), the unit of the U.S. Department of Agriculture that has traditionally been the producer of the nation's soil maps, now recognizes that its future also lies in coordination among the local agencies, farmers, and corporations that have been empowered by technical developments to produce their own detailed data at low cost and in direct response to perceived needs (NRCS, 1995). These arguments for new central roles rest fundamentally on the proposition that only the central government has the power to regulate and advise nationally, and that substantial agglomeration economies still exist in the area of skills and research excellence.

While national government has the advantage of a history of involvement in geographic information production, it is not at all clear why standards should be set at this level, rather than at the local, regional, state, or global levels. The same arguments that support common standards covering adjacent municipal jurisdictions (ease of edgematching, ease of transfer of activities across boundaries, common skills in a larger labor pool) would also support common standards across state boundaries, and national boundaries. Moreover, these arguments are relatively weak: it is not at all clear whether the costs to each of two adjacent states of harmonizing their standards of geographic information production are outweighed by the benefits, and there are no studies of the less tangible social and technical advantages of doing so. How important it will be to maintain institutions at the national level in particular for the purposes of standard-setting remains to be seen.

Recent technical developments are influencing both production and dissemination of geographic information. The key questions of cost-recovery have been discussed elsewhere (Barr & Masser, 1997); here we concentrate on the technical questions of dissemination. In recent years a large number of projects have taken advantage of Internet, WWW, and related technologies to support novel approaches to the dissemination of geographic information that are increasingly replacing the more traditional ones (dissemination of paper, photographic, magnetic, and optical media). These technologies merge several functions: information search, discovery, assessment, and retrieval. They provide analogs to many of the services of the library: collection-building (deciding what information should be loaded into the database, or linked to it); cataloging (building the metadata resources to support information discovery and assessment); circulation (managing the dissemination of information from the database); and preservation (preserving the database contents, archiving historical information). But the culture of the Internet is different in some ways from that of the library. Although both support the notion of open access to information as a public good, the library's control over the admission of information conflicts with the free-for-all Internet philosophy of individual empowerment. The library provides the resources for information abstraction and cataloging of its collection, although these functions are increasingly shared among libraries and with publishers. On the other hand the Internet culture makes the individual provider responsible for metadata production, and must deal with the fact that there are few inducements available to ensure that metadata is produced well. The "gatekeeper" collection-builder of the library paradigm ensures that every entity meets the institution's standards of quality; the anarchy of the Internet culture empowers everyone to be a publisher, and assigns responsibility for assessment of quality to the user.

Yet it is easy to be misled by the technology into believing that these two worlds are fundamentally more different than they actually are. There is no reason why institutional functions of collection-building should not exist in a digital world; it simply takes time for the necessary institutional arrangements to emerge from the apparent chaos engendered by the initial technological change. In the long term, it is in everyone's interests that collections emerge on the Internet whose quality is assured by their institutional sponsors just as in the traditional world. Thus although the library services of collection-building and cataloging may appear to be missing in today's Internet world, the same reasons that led to their existence in traditional libraries will eventually lead to their emergence in the digital world. The Federal Geographic Data Committee's sponsorship of the National Geospatial Data Clearinghouse (www.fgdc.gov) already suggests the pattern of future evolution; information retrieved from this source is subject to the sponsor's standards, and a reputation for bad information will in time undermine the institution. Eventually, users will come to believe in Internet sources of geographic information in much the same way they currently believe in certain map libraries, or publishers.

While much geographic information has a high level of stability, there are good reasons for wanting to ensure that changes are captured and inserted into geographic databases as quickly as possible. Information on changes in demographics on timescales of weeks or months can be of great value to retailers, as they substantially correct the decennial information of the census; current information on road conditions on timescales of minutes can be of great value to drivers. In the traditional world of publishing, information could be maintained by periodically disseminating new versions or updates. Navigational charts provided by the National Ocean Service are reissued at a frequency determined by several factors, including rate of change, and importance of the area to navigation. Between revisions, updates are provided as Notices to Mariners, and regulations prescribe how these are to be maintained and used by captains. In a digital world in which connectivity is assumed perfect and costs of publication have gone effectively to zero, it is possible to devise new arrangements that are more effective. The principle that information should be maintained and served as close to the source as possible (Onsrud & Rushton, 1995) makes good sense, as it gives the originator custodianship, and allows users to "pull" information as needed, with the assurance that the information being "pulled" is as current as possible. Alternatively updates can be "pushed" by the custodian.

However, the connectivity of the Internet makes it possible for a custodian of an IBO to be located almost independently of the IBO for which he/she is responsible. In principle, and problems of bandwidth notwithstanding, it is possible for a local government in California to create, maintain, and disseminate an IBO whose bits are actually located on a server in New York. Why would a custodian of an IBO insist on reducing this geographic separation? Several reasons, all of them fairly weak, appear to be valid:

(1) Customs regulations and other laws argue against locating the bits of an IBO in a different country; tax laws may impede locations in other states.
(2) If the custodian of the IBO is responsible for paying for its maintenance, he/she may well demand that the funds be spent locally. Most jurisdictions give preference to local providers when services are contracted out.
(3) While connectivity may be high, proximal locations may be a better hedge against the possibility of Internet down-time or disruption.
(4) At present, institutions gain substantial prestige from the serving of information; this may be lost if it is known that the information is actually being served from elsewhere.
(5) Serving an IBO from a server physically located within the custodian's institution conveys a sense of control, real or apparent, that may be attractive to the custodian.

INFORMATION OF GEOGRAPHICALLY DETERMINED INTEREST

In central place theory, a good is assumed to be ubiquitous and uniform; the groceries available from one retailer are essentially the same as those available from any other. Practical problems of central facilities location are solved over areas of limited extent, within which there is no substantial variation in the nature of goods. In principle, this same approach could be used at a global level, though it is hard to identify practical examples.

In a system of central facilities dispersing information, such as a system of libraries, the same assumption applies, though to a lesser extent. Variations exist, for example, based on language; the demand for French-language texts is obviously higher in French-speaking countries. The research specialities of a university to some degree drive the contents of its library collection, by providing higher local demand for certain topics.

Variation in demand is much sharper, however, in the case of geographic information. Define a *geographic IBO* (GIBO) as an object containing a representation of variation over some part of the Earth's surface; examples are digital representations of maps or digital Earth images. Define the *footprint* of a geographic IBO as the region of the Earth's surface described by the IBO's contents. Many GIBOs have footprints that are rectangles on some projection. Many IBOs are not geographic according to this definition, but nevertheless have footprints, although the boundaries of these footprints may be less precise, and nonrectangular. Examples of such *geographically related IBOs* include photographs taken at some known location, records from a herbarium that include the location of each collection, guidebooks to cities, or reports on city neighborhoods.

The level of demand for a given geographic or geographically related IBO varies sharply over the Earth's surface. Interest in a digital street map, for example, is much higher in the area covered by the map, and falls almost to zero as distance from the footprint increases. We term such IBOs *information of geographically determined interest* (IGDIs). Of course some degree of geographic variation exists in the level of interest in almost all information. Of interest here, however, are patterns that are well behaved and therefore amenable to modeling; and of sufficient variability to impact locational decisions.

Several models are available to capture geographic variation in demand for IGDIs. A *binary* model would define interest as uniformly high within the extent of the footprint, and zero elsewhere. Let F denote the set defining the footprint, and $p(x, y)$ the probability that an individual located at (x, y) will request the IBO in a given period of time. Then the binary model is:

$$p(x,y) = P \text{ if } (x,y) \in F, \text{ else } 0$$

Alternatively, we might define p as decreasing function f of the distance $d(x, y)$ between the footprint and (x, y) in a *distance–decay* model:

$$p(x,y) = P \text{ if } (x,y) \in F, \text{ else } p(x,y) = f[d(x,y)]$$

The spatial interaction models described earlier provide a reasonable basis for f; interest in an IGDI might be modeled by analogy to interest in some tangible good.

Finally, geographers have frequently made use of hierarchical models of interaction, in which the interaction between two places depends on their levels in some importance hierarchy. Stated in terms of information, this would imply that the level of interest in a given IGDI depends both on the distance between the user and the footprint, and on the user's level in the hierarchy. While the average person in Paris may have no interest in a digital street map of Los Angeles, there will be various individuals and agencies that have some interest, perhaps because of the need to travel to Los Angeles, or because of some information-gathering responsibility. Such models assume that this interest does not extend outside the relevant administrative node to its environs.

The variation in interest in an IGDI runs directly counter to the celebrated expectation that the Internet will "destroy geography"; that in the digital world "there is no more there, everywhere is here". While it is possible in principle to deliver information anywhere, in reality information must be stored somewhere, and there must be interest in it to stimulate delivery. Both storage and interest are thus inherently and persistently geographic, despite the potential for perfect digital interconnectivity. While theories of central facilities location focus on two locations, that of the consumer and that of the offering of the good, the serving of IGDIs involves three: the location of the demand, the location of the server, and the location that the information is *about*.

Consider the problem of searching collections of IGDIs distributed over servers connected to the Internet for information about a given geographic area, or *query footprint*. A query footprint can be defined using coordinates or place names. But in the former case, no current search engine is capable of recognizing geographic coordinates and building catalogs based on them. In the latter case, there is no guarantee that an IGDI covering an area that includes a given named place will actually contain an instance of the name in text, or that the search engine will regard it as sufficiently significant to retrieve and store it. Thus while today's search engines may provide a satisfactory but imperfect solution to the problem of informing the user about IBOs offered on distributed servers, they fail in the case of IGDIs and geographically based search. This conclusion generalizes somewhat; search engines are far less effective than traditional library catalog services at abstracting the key characteristics of an IBO and making them available for search.

Because of this problem, various geographic information *clearinghouses* (GICs) have emerged recently to provide catalogs of geographic information, in effect solving the problem of the user's lack of information about this particular type of IBO, because of the failure of search engines to catalog it automatically. Many state governments now run clearinghouses for GIBOs, and the U.S. Federal Geographic Data Committee sponsors the National Geospatial Data Clearinghouse (www.fgdc.gov). The geographic domain of a GIC is typically defined by the jurisdiction or area of responsibility of the sponsoring agency. With reference to the three locations discussed earlier, a GIC is likely to catalog GIBOs that are about areas within its domain. As argued earlier, because of local empowerment it is increasingly likely that a GIBO is created in or near its footprint. Finally, as argued earlier it is likely that the GIBO is served close to its point of creation. It follows then that a GIC will likely catalog GIBOs that are served from points within or close to its domain.

However, clearinghouses are expensive to operate, due to the costs of creating and maintaining the necessary catalog records or metadata; and rely on information being made available to the clearinghouse by the IBO's custodian. In this sense they are far less effective than the fully automatic search engines. From the user's perspective, a search for a GIBO must follow one of two models: either the user must search all clearinghouses that are deemed likely to include references to the GIBO or the user must search the servers themselves. A digital street map of Goleta, California, may be referenced in the National Geospatial Data Clearinghouse, or in one or more of the clearinghouses maintained by state agencies, or in the clearinghouse of the California Geographic Information

Association, or may be available in a server run by the County of Santa Barbara, or the University of California, or any of a host of other possibilities. This is the digital equivalent of visiting every library, guided by guesswork, but without any form of union catalog. The user will know that because of geographic variation in potential interest, the GIBO is more likely to be found within or near its geographic footprint. But guesswork about the appropriate level of the hierarchy will be much more difficult: is the GIBO more likely to be available in the local public library, or the local university, or the Library of Congress?

TOWARDS A NEW THEORY

This section builds on the previous five to advance a theory of location of information within a digital world, with special attention to geographic information. The discussion is framed within the theory of central facilities location reviewed earlier.

Goods and Services

Information is stored, cataloged, discovered, accessed, and delivered in the form of discrete entities or IBOs, corresponding roughly to today's media-defined books, volumes, maps, images, photographs, and so forth. IBOs are stored on servers, and delivered to user clients for examination, processing, or analysis. Various reasons have been advanced for co-location of many IBOs on a single server; because of a lack of tangible scale economies, servers are expected to be of a wide range of sizes, driven largely by intangible factors such as those discussed earlier.

Consumer Behavior

In the ideal world of perfect Internet connectivity there is no digital equivalent of the impedance of distance, no digital equivalent of range, and demand will be inelastic with respect to distance (although in practice connectivity will certainly continue to be imperfect to some degree). Sharp geographic variation in demand can be expected in the case of IGDIs. The user's desire for one-stop shopping is met by building catalogs of the contents of distributed servers. But while such catalogs can be built very effectively and automatically in the case of textual materials, they fail to support searches based on geographic location, and thus are inadequate for building directories of GIBOs, or more generally for supporting geographically based searches for IGDIs. Instead, agencies will invest in GICs to provide comparatively expensive and limited catalogs of GIBOs within their domains.

Economies of Scale

There are substantial economies of scale to the builders of search engines, because advertising revenue increases with use. No economies of scale have been identified for sponsors of GICs.

Locational Outcomes

The extreme assumptions of classical central place theory lead to regular patterns of offerings on the landscape. When many of these assumptions are relaxed in central facilities location theory much of the geometric regularity disappears, leading to an irregular pattern of offerings that is strongly related to the distribution of demand.

Under the assumptions already discussed regarding the digital world, the relationship between distribution of demand and distribution of supply falls apart. With a search engine providing a catalog, users and IBOs and search engines can be located anywhere. Various arguments have been advanced to support the proposition that IBOs are likely to be located near their originators or custodians.

In the case of IGDIs, however, the locational outcomes are very different. Empowerment of local geographic information production will ensure that originators and custodians are located within or close to a GIBO's footprint, and this conclusion generalizes to IGDIs. An IGDI will therefore also be served from within or close to its footprint. GICs will be built by agencies at various levels of the administrative hierarchy for IGDIs whose footprints intersect their domains. Finally, the user's need for rules to guide discovery of IGDIs and GICs provides additional rationale for local servers. We conclude that while servers of general IBOs may well be footloose in a digital world, the locations of IGDI servers will be driven by strong geographic factors.

CONCLUSIONS

The technological and economic changes being wrought by the introduction of the Internet and other advances in information technology are having two major effects on the production and dissemination of geographic information, both with the potential to impact existing arrangements. By driving the fixed costs of production down, these changes can reverse the traditional arrangements of central production, and empower local entities to produce their own local geographic information to satisfy local needs. Only weak arguments support a continued role for central government, in maintaining standards and clearinghouses, coordinating, and ensuring completeness of coverage.

The future map of research libraries will look very different. Instead of the classical pattern of central service provision that we see today, it will be sufficient for each IBO to be available from only a small number of servers; and under perfect connectivity, from only one. A research library will be able to focus on serving only those IBOs that are of particular relevance to its local role. Its responsibility to a geographically defined constituency argues for it to serve those IGDIs whose footprints overlap its domain, or to provide indirect links to the respective custodians. The library's responsibility to its scholars argues for it to serve the results of their research and their contributions to the corpus of human knowledge, or to provide indirect access to IBOs on each scholar's personal server (though it will likely be argued that the institution is more persistent than the location of the scholar). The library may also serve IBOs that are of particular relevance to the interests of its scholars, or collections of archival IBOs that are analogous to today's special collections. In all other cases, however, the institution will rely not on its own library but on services provided collectively, and paid for collectively. The research library of the digital world will be a much more specialized entity, reflecting the effectively infinite range and zero threshold of library service provision in the digital world.

Note, however, that these arguments derive from consideration of the role of the library as a disseminator of IBOs. Other roles, such as those less prominent roles identified earlier, may well require a pattern of locations much like those in place today. Although the ratio of digital to nondigital IBOs will surely increase in the coming years, it seems certain given the costs of digitizing, and inevitable loss of information that results, that libraries will continue to provide traditional services for many years to come.

If such changes in arrangements are implied by changing economics and technology, then one can legitimately ask how the transition will occur. What steps will be needed to ensure that the transition from central to local production is as painless as possible, and similarly for the transition between research library as central facility, and research library as special collection? In one view, such changes in institutional arrangements are impossible to achieve smoothly, and can only occur through invention of entirely new institutions, and abandonment of old ones. Thus the digital library must be built alongside the existing one, but reflecting entirely new principles of organization and responsibility. This strategy leads inevitably to institutional conflict, as new and old arrangements compete for available resources; and the inevitable decline of the old arrangements can be very painful and wasteful. But smooth transition is possible only if there is consensus within the old institutions of the need for change, and a shared vision of how it can be achieved.

ACKNOWLEDGEMENTS—The support of the National Science Foundation for the Varenius project of the National Center for Geographic Information and Analysis (SBR 96-00465) and for the Alexandria Digital Library (IRI 94-11330) is gratefully acknowledged.

REFERENCES

Barr, R., & Masser, I. (1997). The economic nature of geographic information. In Z. Kemp (Ed.), *Innovations in GIS 4* (pp. 234–248). London: Taylor & Francis.

Berry, B. J. L. (1967). *Geography of market centers and retail distribution.* Englewood Cliffs, NJ: Prentice-Hall.

Berry, B. J. L., Parr, J. B., Epstein, B. J., Ghosh, A., & Smith, R. H. T. (1988). *Market centers and retail location: Theory and applications.* Englewood Cliffs, NJ: Prentice-Hall.

Christaller, W. (1966). *Central places in Southern Germany* (C. W. Baskin, Trans.). Englewood Cliffs, NJ: Prentice-Hall.

Fotheringham, A. S., & O'Kelly, M. E. (1989). *Spatial interaction models: Formulations and applications.* Boston: Kluwer.

Goodchild, M. F. (1972). The trade area of a displaced hexagonal lattice point. *Geographical Analysis, 4,* 105–107.

Goodchild, M. F. (1984). ILACS: A location–allocation model for retail site selection. *Journal of Retailing, 60,* 84–100.

Goodchild, M. F. (1992). Analysis. In R. F. Abler, M. G. Marcus, and J. M. Olson (Eds.), *Geography's inner worlds: Pervasive themes in contemporary American geography.* New Brunswick, NJ: Rutgers University Press.

Hayes, B. (1997, May–June). The infrastructure of the information infrastructure. *American Scientist, 85,* 214–218.

Lösch, A. (1954). *The economics of location* (W. H. Woglom, Trans.). New Haven, CT: Yale University Press.

Love, R. F., Morris, J. G., & Wesolowsky, G. O. (1988). *Facilities location: Models and methods.* New York: North-Holland.

National Research Council. (1993). *Towards a coordinated spatial data infrastructure for the nation.* Washington, DC: National Academies Press.

Natural Resources Conservation Service. (1995). *Data rich and information poor: A report to the Chief of the Natural Resources Conservation Service by the Blue Ribbon Panel on natural resource inventory and performance measurement.* Washington, DC: U.S. Department of Agriculture.

Onsrud, H. J., & Rushton, G. (1995). *Sharing geographic information.* New Brunswick, NJ: Center for Urban Policy Research.

Rushton, G. (1972). Map transformations of point patterns: Central place patterns in areas of variable population density. *Papers of the Regional Science Association, 28,* 111–129.

U.S. Geological Survey. (1997). *National Mapping Division strategic plan.* Washington, DC: U.S. Department of the Interior.