Location Theory: Data Centers & Google

By Niek Goorman (S1396765)
Department of City and Regional Planning,
University of North Carolina at Chapel Hill
Faculty of Spatial Sciences, University of Groningen
Master Thesis, Economic Geography

Abstract

The data center is the beating heart of the modern internet. They serve two important functions: switching data from one network to another, and hosting data (for example, a website). Even though data centers are critical to the functioning of the internet, and even though the internet is nowadays connected to many parts of modern society, there has been very little research about the location of the data center. This paper seeks to explain what factors steer data center location decisions in the United States, and illustrates these factors by examining the location decisions of a new Google data center, currently under construction in Lenoir, North Carolina.

This paper finds that data centers usually locate in urban areas, drawn to the economic and population centers of the United States. However, there appears to be a bias towards those metropolitan areas that serve as 'hubs' for the fiber optic networks in the United states, including those areas where transcontinental cables come ashore. This also holds true on a lower level: data centers tend to locate in those places within cities where multiple fiber optic cables intersect (Evans-Cowley, 2002).

In the case of Google in Lenoir, several important factors were identified: the price of electricity, the cost of land, the availability of fresh water, and the $210 million incentives package provided by state and local governments. Of these factors, electricity prices appear to be the biggest influence, since data centers require very large amounts of power to operate. Electricity prices are very low in North Carolina, and in the Western North Carolina (WNC) area, there was also a large amount of generating capacity available due to the decline of the furniture manufacturing industry in past years.

Table of Contents

List of Figures & Tables

Acknowledgements

I am writing this paper as my Master Thesis, for the Master program in Economic Geography, at the University of Groningen. Part of the research and the writing were completed during my time at the University of North Carolina at Chapel Hill, from August to December, 2007. Therefore, I would like to thank everyone at Neurus, the Network for European and United States Regional and Urban Studies, for having it made possible for me to have spend that semester at UNC.

I would like to thank the students and staff of UNC, and especially professor Harvey Goldstein and Todd BenDor, at the Department of City and Regional Planning, for helping me with my research, pointing me in the right direction and generally being supportive. I would like to extend the same thanks to drs. Paul van Steen, my supervisor at the University of Groningen, for helping me with the administrative side of my stay at UNC, as well as critically examining the draft versions of this thesis, and giving advice about improving them.

Lastly, I'd like to thank all my friends and fellow students in Groningen and Chapel Hill, as well as my family, for all their support and suggestions.

Niek Goorman
Groningen, April 2008.

1. Introduction

The name Google is a name known to virtually everyone who uses the internet. In fact, an online survey found Google to be the most recognized brand among internet users worldwide, easily outranking Youtube, Apple and Wikipedia (Brandchannel 2006). Thus, when Google publicly announced their plans to build a new data center in Lenoir, NC, both online and offline media immediately reported it. However, the same expansion barely drew any attention before the final deal was sealed (local media excepted), when the company behind the expansion was only known as Tapaha Dynamics LLC. Afterwards, Tapaha was found to be a dummy company fully owned by Google.

There are clear signs that Google is rapidly expanding the number of data centers they use. Apart from the Lenoir site, Google is building data centers in many other locations as well, such as South Carolina, Oregon, Oklahoma and Iowa. These sites are anywhere between the planning phase and completion at the time of writing of this paper. The different sites that Google has chosen for their new data centers have one element in common: all are sited in rural areas. This is a relatively new phenomenon: up to now, data centers generally seemed to have been set up in urban areas, as will be shown later on. This paper will examine Google's reasons for expanding into rural Lenoir, NC, while simultaneously providing insight into general data center location theory.

For this project, several research questions were defined that form the subject of this project. The main question of this project is:

- *What are the reasons for Google to locate a new data center in Lenoir, NC?*

Secondary questions are:

- *What are, in general, the location factors for traditional IT businesses?*
- *How do location factors for data centers differ from these traditional IT location factors?*
- *Does Google have additional, specific needs for their data centers?*
- *To what degree are local and regional governments willing and able to meet the location requirements for the establishment of a data center in the Lenoir location?*

- *To what degree can the location requirements of Google's data centers be fulfilled in the Lenoir location?*

The hypothesis concerning these questions is that, today, data centers mostly locate in areas where electricity is cheap, and where large amounts of bandwidth are or can be made available. However, bandwidth is likely of secondary importance as a location factor (at least compared to electricity), because most advanced countries have a well-developed fiber optic network. Even in areas where such a network is underdeveloped or absent, such a network can be laid down or improved either by the data center's owner, or by the state and local governments as an incentive to attract the data center. Several other factors also play roles of differing importance, for example tax breaks from local governments, the availability of labor, and the cost of land. Google likely picked Lenoir because North Carolina has cheap power compared to the U.S. average, and partly because of the existence of a relatively well-developed (or to be developed) fiber optic network in the area. It should be noted, however, that cheap electricity is also available in several other, neighboring states, according to data from the Department of Energy website (Department of Energy 2007). Still, a rural location like Lenoir is relatively cheap to develop, while nearby areas can still provide Google with needed employees, for example through the nearby Catawba College. It is also likely that on the local level, the tax breaks offered by local governments played a role in Google's decision (NCICL, 2007).

When it comes to the general area, the southeast of the U.S., it is likely that Google wants to build their new data center in a location where it can take advantage of the high-bandwidth connections between Atlanta, Georgia, and the U.S. Northeast. This way, Google's services are optimally accessible, much like a company would want to locate alongside a highway connecting major population areas. Possibly, Google also wants to create a network of data centers throughout the United States and the rest of the world, spread out as to optimize both accessibility and security for any location in the world.

For this project, combining and evaluating existing research will be the primary means of finding data. There are papers written on data center requirements, on the location of IT-related activities in rural areas, and there is some information available on Google's data centers online (see above). However, to date there has been no academic effort made to get the bigger picture as to why Google and other companies are locating data centers in specific locations.

Google’s Lenoir data center will be looked at to determine Google’s requirements for their decision to locate a data center there. Other locations where Google has a data center, or where one is being planned, may be used to provide a stronger case for the importance that Google places on any particular location factor. Because of the strong focus on Google's Lenoir data center, this project shares some similarities with a case study. However, it does not meet the standards set in Hancock & Algozzine (2006), who characterize a case study as research meeting three requirements:

- Usually focuses on a phenomenon (an event, program, situation or activity).

While this research could be seen as an event (the location decision), it is not focused on the consequences or wider impact of the event, but more on the event itself, and how it came to be (why the location decision was made).

- Aimed at the event or phenomenon in its own space and time.

It is important to look at Google's decision in the right context: there is a reason why IT companies have only recently started to move to rural America. However, there are few social and cultural elements to be explored here: the context is not as important as it might be in other research.

- Case study research is richly descriptive.

The research will be exploratory in nature; providing insight in location factors of data centers, focusing on Google's, will be its goal. While description plays a role, it is not the primary goal of this research.

The method used in this research will be partly quantitative and partly qualitative. Quantitative data can be found on topics like electricity prices (for example, through the website of the Department of Energy), land costs, and labor costs/cost of living in Lenoir. These data are focused on reviewing the location of the data center. Qualitatively, the project will focus on 'getting the bigger picture', through reviewing literature and conducting interviews. Here, the main focus is on the reasons behind data center expansion to certain, in this case rural, areas, and Google's need for these data centers. More specifically, one of Google's reasons to locate in the western NC area may be to form a wide-spread network of data centers around the US. This cannot be easily measured in a quantitative way.

Obtaining original data proved to be difficult, as there is much secrecy surrounding telecommunication companies' inner workings and infrastructure, something that Moss & Townsend also noted in their 1998 paper. This secrecy was stepped up after the 2001 attacks on the World Trade Center in New York. In 2003, the U.S. Department of Homeland Security became aware of a George Mason University student who was conducting research into the vulnerabilities of infrastructure in the U.S., among which was the fiber optic network. While the research was ultimately published, some more sensitive data was left out, and some sources on U.S. infrastructure (concerning communications and utilities) were removed from the public domain as a result (Washington Post, 2003).

Governments, however, do not have a monopoly on secrecy. Google is notorious for releasing very little information regarding their operations. For example, they have not publicly released figures on the number, size, or locations of their data centers. Some estimates are available on various websites, however (threadwatch.org 2006, Baseline Magazine 2006), though how correct these numbers are will need to be determined. One way to obtain some 'original' data, however, is through an interviews, conducted by email or telephone. Unfortunately, distance and transportation problems made it impossible to conduct face-to-face interviews with relevant actors, but emails and phone calls have still proven helpful to this research.

This first chapter has served as an introduction to this paper. The second chapter will serve as an overview of the location theory of IT businesses, to provide some contrast with the location preferences for data centers. Where IT firms tend to locate is the focus in this chapter. Following that, the third chapter will provide an overview of the history of the internet, from the creation of ARPAnet to the complex network it has become today. The fourth chapter will give a short history of Google, both as a company and as part of the internet’s infrastructure, and will give some insight into the possible number and locations of Google’s data centers. This chapter will also seek to explain the normal operations of a data center, as well as the spatial and other requirements a data center needs. Finally, it will provide a short overview of the similarities and differences between location theory for general IT businesses and data centers. Chapter five will then give an overview of the Lenoir area, where Google is locating their latest data center, and will include a short overview of the deal Google brokered with the various actors in the area, at least that what is publicly known about it. Lastly, chapter six will serve as the conclusion to the paper, and the research questions posed in the introduction will be answered there.

2. Location Theory: Traditional IT Businesses

One of the founders of location theory is Alfred Weber. One of his most important works is the *Theory of the Location of Industries*, written in 1929. It was the first work to cover the location preferences of the manufacturing industry, something that was new to the academic world. Up to that point, as Weber himself explains, economic geography was only concerned with the spatial distribution of industry in a purely statistical way, creating an index of what regions were home to what types of industry. Weber built his ideas on those of Von Thunen, who created a general location theory in the late 18th century that focused on the existence of different land uses in concentric rings around cities, based purely on the cost of land and the distance to the market (the city center). Weber, however, looked at what drew certain types of industry to a certain area, and he explained why industry favored certain locations over others, apart from cost of land and distance from the market. He then classified the location factor he found into several groups. For example, Weber distinguishes between general factors, like rent, transportation and labor costs, and special factors. These are factors that influence only a particular industry, like the local climate, or the availability of a source of fresh water (Weber 1929/1957).

The above applies to the manufacturing industry, of which the location factors are now well understood. Compared to manufacturing, the study of location factors of the service and IT/communications sectors is relatively new, and their spatial behavior is not yet as well understood. The service sector in itself is huge, and as such it is hard to pinpoint particular location factors. Strongly connected to location factors, however, is the concept of clustering. Clusters were first described by Marshall, in 1890, and were more recently popularized by Porter in 1990, leading to what is now called New Economic Geography.

Clusters are one location factor that affects many industries, both in the manufacturing and the services sector. Essentially it describes a group of firms that have located in proximity of each other. Apart from the general properties of an area, leading to certain costs and benefits (called geographical benefits), clustering leads to agglomeration benefits. Some examples described by Maggioni (2002) are lower costs of doing business, due to the ability to negotiate with suppliers together, or due to economies of scale making it cheaper and easier for local governments to provide adequate infrastructure in the area. Another agglomeration benefit is the 'connectedness' that a firm has when it is located nearby other, similar firms. Knowledge flows easier between firms located closer to one another (also called information spillover), and a cluster will also lead to better qualified employees available in an area. All these advantages can make the presence of a cluster a desirable location factor for a firm.

For the Seattle metropolitan area, research into location factors of high-tech firms has been conducted by Sommers et al. (2000), who interviewed a number of firm executives, real estate managers, and municipal planners about location preferences for high-tech firms. From this, Sommers et al. made several 'fairly consistent' observations concerning the locations of certain types of high-tech business. First of all, high-tech manufacturing in Seattle generally prefer suburban locations, because of current and future space requirements that can not be fulfilled in downtown areas. Second, e-commerce, internet content and telecommunications companies prefer the inner city. Software companies are less specific in their location preference, and tended to locate in and around redeveloped areas close to downtown. Lastly, biotech firms are attracted to hospitals and research institutions like universities, possibly as spinoffs, though any manufacturing facilities are generally placed on suburban industrial parks. In general, these firms did not move between urban and suburban locations very often, and generally remained in their own area (Sommers et al. 2000).

One interesting observation made by Sommers et al., is that of the influence of the employees on the location of a business. They find that the founder of a firm often locates it in a place with a high quality of life, and refer to research among biotechnology firms showing that the residential location preference of a firms' founder or CEO is the single most important factor in locating these type of firms (Sommers et al. 2000). The role of personal preference of employees and high-level executives should not be underestimated, especially in smaller, more specialized firms that are not bound to a certain area. Software developers are an examples of these types of companies, that are generally referred to as 'footloose' companies, because they do not have pressing reasons (like a needed resource or close proximity to a large customer base) to permanently reside in any given area.

Concerning software companies, IT-firms are sometimes regarded as an opportunity to revitalize otherwise lagging economic areas (Camagni & Capello 2005). However, Camagni & Capello also warn against the image of IT-related businesses as a foolproof way to economic renewal in an area. Still, Milligan (2001) shows that, for the majority of surveyed local governments in the United States, the technology & telecommunications sector is one of the main current focuses of their local economic development program, second only to the manufacturing sector. Table 1 shows a comprehensive list of figures regarding the expected future economic base of local governments in the United States. In this 1999 survey, 11,5% of all local governments (including cities, towns, villages, counties and boroughs of all sizes) reported that the telecommunications & technology sector would constitute this base for the period 1999-2004. Table 1 also shows that the telecommunications & technology sector was expected to develop in cities and towns of all sizes, even though there is a tendency towards bigger cities[1].

Population Group	Respon -dents:	Agri-culture	Manu-facturing	Retail/ service	Insti	Res comm	Tour/ hosp	Ware/ dist	Tech/ tele	Other
Total	889	3.8%	19.7%	26.7%	7.4%	16.2%	5.4%	4.5%	11.5%	4.8%
Over 1,000,000	7	0%	0%	28.6%	0%	0%	0%	14.3%	42.9%	14.3%
500,000-1,000,000	12	0%	16.7%	33.3%	0%	0%	0%	0%	25%	25%
250,000-499,999	32	3.1%	18.8%	25%	9.4%	9.4%	12.5%	6.3%	12.5%	3.1%
100,000-249,999	92	4.4%	19.6%	29.4%	7.6%	4.4%	9.8%	0%	19.6%	5.4%
50,000-99,999	142	3.5%	26.8%	32.4%	6.3%	9.2%	2.8%	4.9%	10.6%	3.5%
25,000-49,999	207	1.5%	12.6%	33.3%	9.2%	15.9%	4.8%	5.3%	11.6%	5.8%
10,000-24,999	397	5.3%	21.4%	20.4%	7.1%	22.9%	5.3%	4.8%	8.8%	4.0%

Table 1: Expected economic base of responding local governments, 1999-2004, during survey held early 1999. Insti=Institutional. Res comm= Residential Community (commuters). Tour/hosp=Tourism & Hospitality. Ware/dist= Warehousing & Distribution. Tech/tele=Technology & Telecommunications. Source: Milligan (2001).

Another example of research into the spatial spread of IT firms has been conducted by Atzema (2001) in his analysis of the location of almost 19,000 IT-related firms in the Netherlands. Unsurprisingly, he concludes that most IT firms can be found in larger cities, with a minority locating in smaller cities or rural areas. The majority of IT firms in the Netherlands are located in or near the western part of the country, most likely because this area is both the economic heart as well as the most densely populated area of the Netherlands. However, Atzema also cites location factors of IT firms, on the basis of factors cited by the firms themselves in a survey (see figure 1). From this, it appears that IT firms find it very important to be located in places with optimal accessibility by car. While this seems surprising for a type of business related to intangible things like software

[1] Note that the sample size for the two categories with the highest population in table 1 is very low; decreasing the reliability of those particular results.

and telecommunications, it can be explained, as indeed Atzema does, by the fact that many IT firms rely on consultants and technicians who travel around the country, going from customer to customer to provide support for the firms' product, and thus a central location decreases their average travel time. However, Atzema also distinguishes between large and small firms, noting that the larger firms, dominating the national market on IT, are more interested in a central location, as to remain optimally connected to all parts of that market. Smaller firms, in comparison, operate on the local or regional level more often, and are thus more interested in regional accessibility and a location that is only central on the regional or local level (Atzema 2001).

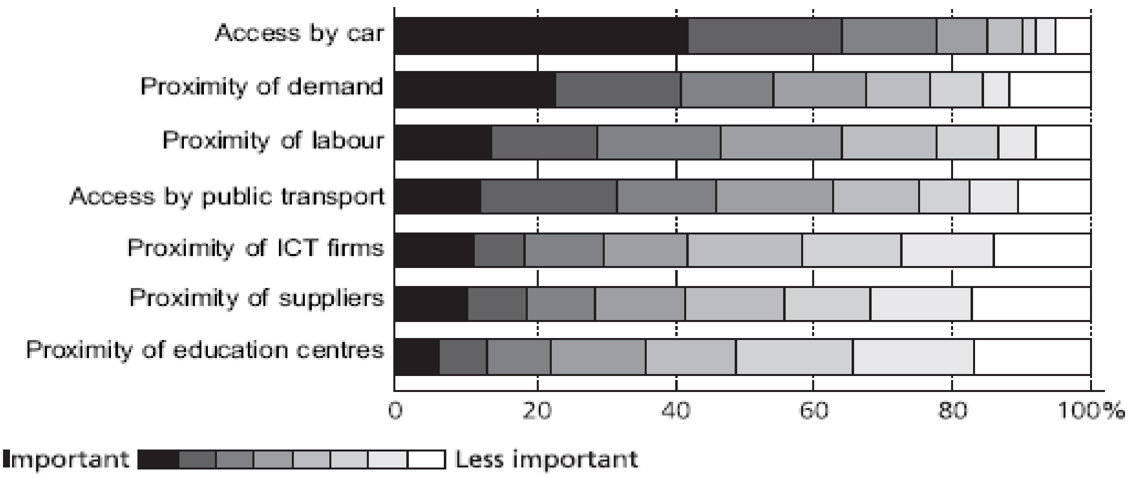

Figure 1. Survey results on location factors of IT firms in the Netherlands. Source: Atzema (2001).

Concerning the United States, the largest concentrations of IT businesses exist in the San Francisco metropolitan area and the northeast, specifically in and around Boston. Cortright & Mayer (2001) conducted a study on the general spread of high-technology businesses among 14 metropolitan areas in the United States, and concluded that metropolitan areas specialize in only one or a few fields of high-tech expertise. The only exception Cortright & Mayer found was Silicon Valley/San Jose, where many different fields of expertise are present. The metropolitan areas with the most jobs in the field of information technology are Boston, Washington D.C. And San Jose[2]. Of these three cities, Washington DC specializes in data processing, and is not as important as Boston and San Jose on the issue of innovative IT businesses. Cortright & Mayer also take the degree of specialization into account, showing that information technology also accounts for a large part of the local economies of Portland, Seattle, Austin, Sacramento and Denver, even though the absolute number of jobs in IT is lower here than it is in Boston, Washington D.C. and San Jose (Cortright & Mayer 2001).

It is difficult to point out any specific location factors for general IT businesses in the United States. Mostly, it seems they are simply a product of path-dependency: the IT sector in an area flourishes because the area has traditionally housed high-tech, innovative industries. For the Boston and San Jose/Silicon Valley areas, this has been documented by Hulsink, Manuel & Bouwman (2007). In the case of both these areas, high-tech industries were active in the area since at least the 1950s, and were often (linked to) defense contractors.

[2] The definition of information technology used is the categories 'software publishing' and 'information services & data processing services' from Cortright & Mayer (2001), table 2. Unfortunately, there is no international standard for the term, and thus Cortright & Mayer's definition is different from the one used by Atzema.

In conclusion, it appears that traditional IT businesses value their connections to their respective markets, as well as accessibility by car. As Atzema (2001) shows, the exact preference is likely to be related to the size and focus of the business in question, with larger businesses attempting to connect to the market on a higher level than smaller businesses, which prefer local or regional accessibility. This accessibility translates into choosing locations central to the markets served, as well as physical accessibility by car, or sometimes public transportation (Atzema 2001, see figure 1). Part of the location decision of an IT firm is historical, according to Hulsink, Manuel & Bouwman (2007), who show that the two largest IT clusters in the United States have evolved from earlier high-tech clusters that were often linked to defense contractors.

When it comes to small, newly started firms, Sommers et al. (2000) say these have their location decision dictated, or at least influenced, by the founder or other personnel employed by the firm. At the same time, Milligan (2001) shows that a majority of local and regional governments in the United States have made the telecommunications sector their prime concern when it comes to attracting new businesses, and table 1 proves that technology and telecommunications was expected to become the economic base of not only large cities, but also of some small towns, coinciding with Atzema's (2001) idea of some IT businesses aiming specifically at a local or regional market.

3. The Internet: History & Structure

Originally created by the United States military during the Cold War, the internet's predecessor was constructed to ensure that military communications would remain functional, even after one or more systems would have been disabled, for example by a Soviet nuclear attack. Essentially, it was created to decentralize the communications infrastructure so as to make it much more reliable in case of failure of one or more of its components. This network was called ARPAnet, after the Advanced Research Projects Agency that commissioned it. Later on, as Hillstrom (2005) explains, this proto-internet started to be used for non-military purposes, like connecting universities and other research institutions (like defense contractors) with each other. The network became active first in 1969 with only four nodes connected (a node is an independent system on the network). In later years, ARPAnet grew rapidly to 15 nodes in 1971, to 200 in 1981, to 2000 in 1985.

Further internet history is explained well by Abbate (1999), in *Inventing the Internet*. In the early 1980s, the U.S. military created its own network, called MILnet, separate from ARPAnet. This freed up ARPAnet for more non-classified purposes, and led to the creation of a network connecting several computer centers at various universities, aimed at linking together, as well as providing easy access to, computing resources for research purposes. This, in turn, led to the creation of a dedicated academic network called CSnet in 1983, which was later superseded by the more well-known NSFnet, named after the National Science Foundation which funded the project. This network included high-speed, mainline fiber optic cables (called *backbones*), and thus was the main network provider for internet access until the early 1990s. Many more networks existed (for varying amounts of time) during the 1970-1995 period (for example, Telenet, BITnet, and Phonenet). These networks were usually regional, connected with each other using several different internetworking technologies, but in the end it was the TCP/IP protocol that won out, and is still used to this day[3]. Thus, by linking the different computer networks, the internet was created, although access was still more or less restricted, and it was NSFnet (who operated the main backbones at the time) deciding who could and could not get access (Abbate 1999).

Ultimately, ARPAnet was decommissioned in 1990, its remaining users being transferred to the NSFnet. At this point, NSFnet was slowly opened to the public, and the parties investing in the network shifted from governments and non-profit institutions to the private sector (Hillstrom 2005). This spawned unprecedented growth of the network, both in terms of data traffic, the amount of

[3] For a more detailed history of the TCP/IP protocol, the early ARPAnet, and its creators, see Hafner and Lyon (1996).

connected nodes, as well as the actual, physical infrastructure needed to allow for this growth. The internet, at this point, was almost fully overlapping with NSFnet, and the two became largely interchangeable. The National Science Foundation still controlled access to the network, however, including their Acceptable Use Policy, which stated that the network was only to be used for research and non-commercial activities. Pressure from commercial parties to gain access to the network, as well as to set up competing backbone infrastructure, eventually led to the NSF slowly transferring responsibility of the network to various other parties. Various commercially-operated backbones were put into place in the early to mid 1990s, and NSFnet was formally decommissioned in 1995 (Abbate 1999).

Upon the decommissioning of NSFnet and the move to the current form of the internet, four Network Access Points (NAPs) were created. Essentially, these NAPs served to connect existing regional networks with each other, a task that was, at that point, still handled by 16 NSFnet nodes. The NAPs were owned by the National Science Foundation, although the operations were contracted to private firms, as opposed to the non-commercial, publicly operated NSFnet. This switch happened through several steps. First, NSFnet connected to the NAPs. Then, regional networks connected to the NAPs through Internet Service Providers (ISPs). Lastly, these regional networks were disconnected from the NSFnet, so that access to the old NSFnet backbone now ran through ISPs and the four Network Access Points (Harris & Gerich 1996). At this point, the NSFnet architecture was decommissioned, and the internet took the shape it still has today: high-bandwidth fiber optic cables connecting access points with each other, and ISPs connecting their customers to these access points.

However, the huge growth in data traffic over the years meant that having only four points where the different networks connect to each other became very inefficient. The four NAPs became heavily congested, and so private companies began to invest in the internet's fundamental infrastructure. As commercial use of the internet grew, network providers started to lay down new high-capacity fiber optic backbones. The market also started connecting both the new and the existing networks with each other through new data centers situated in places where these networks intersect, allowing for a transfer of data from one network to another, providing the same service that the NAPs did. The NAPs were essentially data centers: buildings where multiple networks, made up from fiber optic cables, converge. Here, internet traffic is exchanged between the networks, so that each computer on one network can communicate with any other computer on any other linked network. Today, most data centers offer various other services instead of just connecting different networks (a service called *peering*), but the connections from backbone to backbone remain vital to the internet's functioning. More detailed information on peering and access points (public and private) can be found in a paper by Malecki (2002).

In this chapter, a brief overview of the history of the internet has been given, as to provide a general image of the internet's infrastructure that is still in place today. The internet originated as a project from the U.S. Department of Defense, to allow for reliable communication across a communications network, even when part of the network would be disabled. The current internet basically consists of many smaller networks linked together, so that every computer on one network can communicate with every other computer on every other network (Abbate 1999). Linking these networks together is a major part of why data centers exist, because they allow data to be transferred from one network to another. While the internet's infrastructure used to be publicly operated and maintained, it is now largely privatized, including the backbone networks, and access is no longer regulated through a single institution.

4. Data Centers & Google: History, Definitions & Locations

In 1998, the internet was already widespread in most modern countries. 1998 was also the year that Larry Page and Sergey Brin laid the foundations of their later success, by commercializing their innovative search engine technology into a company called Google. As explained by David Vise and Mark Malseed in *The Google Story* (2005), Page and Brin initially tried to offer their search technology to various existing companies, including competitors-to-be Altavista and Yahoo, but their offers were turned down time and again. Only then did Page and Brin decide to quit their PhD program at Stanford University and founded Google. The story of Google is likely the most successful story of any 'New Economy' company still in existence. At first, they were similar to any other obscure internet startup. In 1999, Google, Inc. first moved into actual offices in Palo Alto, California. At that point, their search engine already handled 100,000 search requests per day, rising to 500,000 later that year (Vise & Malseed 2005). The amount of daily search queries through Google's websites continued to expand rapidly in later years. In December 2006, Google's websites already handled 4,3 billion search requests that month, or over 143 million per day. One year later, in December 2007, the numbers were 30% higher, with 5,6 billion search queries that month, or over 186 million per day. With the total amount of internet search queries measuring over 9,6 billion per month, this puts Google at a market share of about 56% among major search engines. Since the total number of searches only increased by 15% in 2007, this means that Google's market share also grew significantly that year (ComScore 2008).

While search engine technology was Google's great success, they have since expanded their business to email, online photo albums, satellite imagery, and much more, as well as numerous business-level services. Essentially, all of Google's services are in some way related to making it easier for people to access information. Recently, however, Google's semi-philanthropic arm (called Google.org) announced plans to start investing in renewable energy (Google 2007)[4].
To power Google's internet services, huge amounts of bandwidth, processing power and storage space are needed. After all, their many services have attracted millions of users, seeing as these services are largely provided for free, though supported by advertisements. The resources needed to keep these services running reliably can only be provided by large-scale data centers. However, the exact number of Google's data centers is a closely guarded secret. Online estimates seem to reach consensus that Google has about 45 to 60 data centers worldwide (threadwatch.org 2006, Baseline

[4] Google also has an interest of their own here: one of the goals of the project is to make renewable energy cheaper than coal, which would also lower Google's (tremendous) electricity expenses.

2006). However, it is hard to say how reliable these figures are, and the numbers may well be out of date. A more recent article on datacenterknowledge.com (2008) puts the number at 'at least 12' major data centers in the U.S., and several in Europe.

First, there is a need to define what exactly a data center is. A data center goes by many different names, and, simultaneously, there are several sorts of telecommunication facilities that are sometimes all called 'data center', (for example, they are also called telecom hotel, data hotel, hosting facility or carrier facility). Evans-Cowley (2002) distinguishes three types of telecommunications facilities. One type is where the owner runs and utilizes the facility entirely, and owns and maintains the equipment themselves (a carrier-owned co-location facility). A second type is the carrier-neutral co-location facility, where companies lease just the space and utility hookups from the facility owner, but bring in their own equipment. The third type is the data center, where a facility is owned and run by a single company, and customers can rent space, bandwidth and other services on equipment owned by the facility owner (Evans-Cowley 2002). However, the difference between the three types of facility is often minimal, and mostly concerns ownership of the facility and equipment. Therefore, all three types of facility will be referred to as a 'data center' from here on, in accordance with the more general use of the word.

Evans-Cowley (2002), in *Telecom Hotels: A Planners Guide*, defines a data center as 'a shell building with utility hookups and connections to the fiber-optic information superhighway' (ibid. p.1). This immediately defines the most important needs of a data center: access to telecommunications infrastructure, access to (other) utilities to enable this function, and a building to house the necessary equipment. For most data centers, telecommunications infrastructure means having (access to) a dedicated fiber optic cable connecting the data center to the internet. Telephone lines or a satellite connection could also fall in this category, but these are not relevant to the average, internet-based, data center under discussion here. When it comes to the other utilities that enable a data center's operations, electricity is the single most important one. Without it, a data center can not function, as it is fully reliant on computer equipment. For this reason, many (if not all) data centers have their own power generators to provide backup electricity in case of power outages in the area.

Evans-Cowley goes into a fair amount of detail on where data centers are located in cities, and cites actual examples of data centers present in several urban areas, as well as examples of ordinances aimed at attracting or steering data center growth within the city. Evans-Cowley found that data centers would often locate in old, refitted industrial buildings, because these buildings tend to meet certain space and layout requirements. For example, a building that offers a large amount of floor space, with reinforced floors and high ceilings would be advantageous, as it allows for the presence of bulky, heavy computer equipment, while still allowing for the necessary air flow and thus cooling. Another property of these industrial buildings is that they often have a dedicated loading dock that data centers can use to load and unload computer equipment from trucks. The ideal locations of these buildings would be near, next to, or on top of as many backbone networks as possible, to provide links between the different networks, which in turn will ensure maximum connectivity from multiple areas served by those networks (Evans-Cowley 2002).

There are more reasons for location preferences for data centers, however. A data center is, after all, a building providing a certain service, and as such there are places where such a building is easier to build, where the service is easier to sell, where the equipment is easier to maintain, where labor is cheaper or of higher quality, et cetera. Early works on data centers already outlined some of these points, for example Halper (1985). For one, Halper considered electricity to be an important factor, but expected the importance of electricity as a location factor to diminish in the future. While his prediction initially came true, electricity is now one of the most, if not *the* most, important factors in data center location due to a general rise in the cost of energy. Department of Energy (2007) statistics show a large increase in energy cost in the past years, and there are large differences in the cost of electricity between one region of the U.S. and another (Department of Energy, 2006). Regions with low electricity prices will attract data centers to some degree, as can be seen in Google's decision to locate in Lenoir and several other sites with favorable electricity rates..

A second factor Halper lists is telecommunication resources (telephone lines and fiber optic cables). Halper stresses the importance of situating the data center in an area where power and communications cables are available below ground, as to reduce the chances of interruption of service due to weather, accidents or sabotage (Halper 1985). Highway access is another factor mentioned by Halper, though the reason he states for this, ease of transporting data to a backup site, has become redundant due to advances in telecommunications. The same argument was made by Schaeffer (1981), and he also mentions the other arguments Halper makes. Other factors mentioned by Schaeffer are low rental and tax rates, a professional police and fire service (as opposed to a volunteer one), and a good nearby education system to attract or train qualified employees.

Data centers not only connect different parts of the internet with each other, they often serve (exclusively) as data hosting facilities, serving web pages and files to internet users. Because of the nature of data centers (facilities full of computer equipment, operating 24 hours a day, 7 days a week), as well as the increasing dependence on computer systems in the last decades, electricity is now one of the most important factors concerning data center location. A recent case study among 22 data centers in California and New York shows that the average data center uses 52 watt of power per square foot (Koomey 2007). This means that a 100,000 square feet data center would use as much as 5,2 Mw of continuous power. Per month (30 days), electricity usage would be 5,2*24*30= 3744mWh. The average North Carolina household, for example, consumes 1098kWh per month on average (Department of Energy 2007). This means that a 100,000 square foot building[5] at the Lenoir site would use the same amount of electricity as about 3400 households. The cost of industrial electricity in North Carolina is currently 5,23 cents per kWh, which means that the monthly electricity cost of said facility would be almost $200,000. Since Google is building two facilities on their site in Lenoir, this would mean a monthly electricity usage of 7488mWh at almost $400,000 per month, though this excludes any bulk discounts that Google has likely been able to negotiate with Duke Energy, the local electricity company that helped Google set up in the area (see chapter 5).

However, it should be noted that Koomey's 2007 estimates are based on an average amount of electricity among a relatively small sample of data centers in California and New York. In contrast, datacenterknowledge.com (2008) estimates Google's total electricity usage for any of their major data centers at 'at least 50 megawatts', and quotes even higher estimates. However, this figure contrasts with the above in that it covers the entire site, not just one building. Still, a 50Mw figure for the entire site at Lenoir would put the monthly usage at 36,000mWh, or as much as over 30,000 average North Carolinian households, at a monthly cost of (barring bulk discounts) almost $1,9 million.

As of the location of Google's data centers, little is known about them, other than where new ones are currently under construction in the United States. However, there are lists of IP-adresses in use by Google that circulate on the internet. An IP-adress is a unique code, between 4 and 12 digits, that identifies a device connected to the internet or a local network (for example, the University of

[5] It should be noted that, while the 100,000 square foot figure is mentioned in various articles, Google spokesman Barry Schnitt would confirm nor deny it, according to another article on datacenterknowledge.com (2007). However, the Google data center under construction in The Dalles, Oregon has been determined to be 100,000 square foot (Harris 2007), and this is most likely correct for the Lenoir site as well.

Groningen website's IP-adress currently is 129.125.50.68). However, the IP addresses used by Google eventually refer to web addresses in the format gfe-xx.google.com, where the xx stands for a two letter code (sometimes three) that seems to refer to a unique data center, or part of one. For example, http://gfe-ed.google.com is a valid address that will refer a browser to the standard Google search engine, except that it refers to the version of Google's search engine hosted in a particular location. These versions should be almost exactly the same, but small differences may exist. Normally, when simply going to http://www.google.com, your computer's request will be referred to the data center that can process your request the most efficiently. Usually, this will be the data center that is closest to you, but it may also be a different one, should the nearest one be unreachable or congested.

For any web or IP address, it is possible to run a so-called traceroute on it. This can be done from Windows itself, and will list all the nodes on the internet that a signal needs to pass between the user's computer and said web or IP address. For example, doing a traceroute on http://www.google.com, will first show the signal passing from the user's location to the Internet Service Provider used, after which the signal is likely transferred to a backbone network. From there, the signal makes its way to Google's own internal network, likely at their nearest data center. The longer the signal needs to travel to arrive at its destination server, the more nodes it will usually pass. From the information returned from each node, it is often possible to map its location. A program called VisualRoute was used to find more detailed location of each possible gfe-xx.google.com location (GFE is thought to be an abbreviation for Google Front End). The list containing the gfe-xx's were taken from several websites (Threadwatch 2006, Baseline 2006, SEOP 2008).

The results are not very clear. Of the 45 GFE's known at the time of writing this paper, 21 end in Atlanta, Georgia, with the other traceroutes ending in Colorado, Virginia, Washington DC, Chicago, Amsterdam and the United Kingdom (see table 2). Many of the traceroutes end in Atlanta, Georgia. Since this is a major hub for internet traffic, a place where many backbone networks intersect, it is possible that Google has several data centers here because it is such a central location. However, the sheer amount of GFE's ending here may also indicate that this is simply the point from where the trace route can no longer be followed or identified, so it merely seems that the connection stops there. It is also possible that from here, the traffic is routed to other Google data centers through fiber optic cables owned by Google.

Location:	GFE traceroutes ending there:
Atlanta, Georgia	21
Broomfield, Colorado	8
Ashburn, Virginia	5
Amsterdam, Netherlands	3
United Kingdom	3
Hemdon, Virginia	2
Washington, DC	2
Chicago, Illinois	1

Table 2: Possible numbers and locations of Google data centers. Source: Baseline 2006, Threadwatch.org 2006, author.

When it comes to the spatial preferences of data centers, two factors can be identified that coincide, fairly accurately, with where data centers are generally located. These factors are population and connectivity. First, it can be clearly seen that many current data centers are located in or near urban centers in the United States. On datacentermap.com, a website dedicated to mapping the geographical location of data centers in the U.S.[6], the three largest cities in the United States (New York, Los Angeles, and Chicago), all have a large amount of data centers present in and around them. Other large cities show similar patterns. This pattern, however, is altered by connectivity. Some cities are more well-connected to the internet than others, and this, too, can be seen in the spread of data centers in the U.S. The Denver metropolitan area, for example, only has about 2,4 million residents, yet has 9 listed data centers. By comparison, Houston only had 8 data centers listed, even though its metropolitan area contains over twice as many residents. The reason is simply connectivity: Denver is located centrally in the US. When the internet was first created, major cities were slowly connected with each other, until it reached from the west to the east coast. Naturally, the largest cities in the middle of the United States became hubs connecting the coasts, and this situation continues to this day. Houston, in comparison, is fairly well-connected because of its size and regional hub function, but does not connect major areas of the United States the way cities like Denver or Atlanta can. This connectivity is not just visible on the datacentermap.com map, but also on the network maps of the larger network providers, like Level3 (2006) and Cogent (2008). Both network providers use the Denver metropolitan area as their main hub. This can also be seen on a 2000 map of the internet's infrastructure (see figure 2), listing 41 backbone networks, before commercial and security interests, as well as increased complexity of the infrastructure, stopped further public comprehensive mapping of the U.S. fiber optic network.

[6] There are, however, several limitations to the data on datacentermap.com. First of all, for any data center to appear on the map, it must be submitted to the website by a user, so more well-known or more highly used data centers are more likely to appear on the map. Second, the website only lists data centers actually selling capacity or services to the general public. Dedicated data centers fully owned and used by a single company, like Google, are not listed. However, at the time of writing the website lists 339 data centers in the US and is likely to give a fairly accurate image of the spread of data centers in the United States.

Figure 2: 'Largest combined intermetropolitan links on 41 backbone networks, 2000'. Source: Malecki 2002.

When the spatial preferences of general IT businesses are compared to those of data centers, there are both similarities and differences. The IT industry, as explained in chapter 2, tends to locate central to their markets, which are usually major metropolitan areas and economic centers. In the United States, the most important IT clusters are found in Boston and San Jose, and these came to be through historical developments (Atzema 2001, Cortright & Mayer 2001). The similarity with data centers is that they, too, tend to locate near major metropolitan areas and economic centers. There, they prefer locations where multiple fiber optic cables intersect. The major difference with general IT businesses, however, is that this preference can also be seen on a larger scale. Data centers are drawn to those metropolitan areas where they are optimally connected to the internet. From the map at datacentermap.com (2008), combined with maps of the U.S. fiber optic network (see figure 2), it appears that data centers are abundant in and around cities like Denver and Miami, more than could be reasonably expected there when looking at the population size and economic importance of those areas. Denver is the major hub for fiber optic backbones connecting the two coast of the United States. Miami is one of the cities where many major transatlantic backbones come ashore.

This chapter attempted to explain what a data center is, as well as where they tend to locate. Evans-Cowley (2002, p.1) defines a data center as ‘a shell building with utility hookups and connections to the fiber-optic information superhighway’. Often, they house the computer equipment that host websites, making them accessible through the internet. The most critical needs of a data center are electricity and communication links, usually fiber optic cables. A data center tends to locate among intersections of fiber optic cables. This can be seen both on the city level (Evans-Cowley 2002) as well as the national level, where data centers appear to have a bigger presence around those cities that serve as hubs for major fiber optic cables. Another focus of this chapter was Google. It appears that Google’s need for additional data centers can be explained by looking at the tremendous growth of the company in past years. However, the company has never publicly stated the number or location of their data centers, and estimates vary between roughly 20 and 60 worldwide. While it appears that Google has spread out their operations throughout the U.S. and the world, the specific locations of their data centers remains unknown, although data indicates that there may be clusters of Google data centers operating in and around major fiber optic hubs.

5. Lenoir: Regional Overview and the Google Deal

The site that Google has ultimately picked as the location for their new data center is located in the city of Lenoir, in Caldwell County, in the western part of North Carolina (see figure 3). The city of Lenoir is the county seat, with a total county population of 79,509 (SRC 2007, using 2006 data). The entire area is largely rural, with the closest major city being Charlotte (population 664,342), 70 miles to the southeast. According to data from the U.S. Census, Caldwell County's number of elderly people is higher than state and national averages at 14%. Racially speaking, the county is very homogeneous with a 93,3% white population. The average household income in Caldwell is, at $36,748, lower than both state ($40,863) and national ($44,334) averages. However, despite the relatively low average household income, the number of persons below the poverty line is slightly lower than the state average, although it is still higher than the national average US Census 2007).

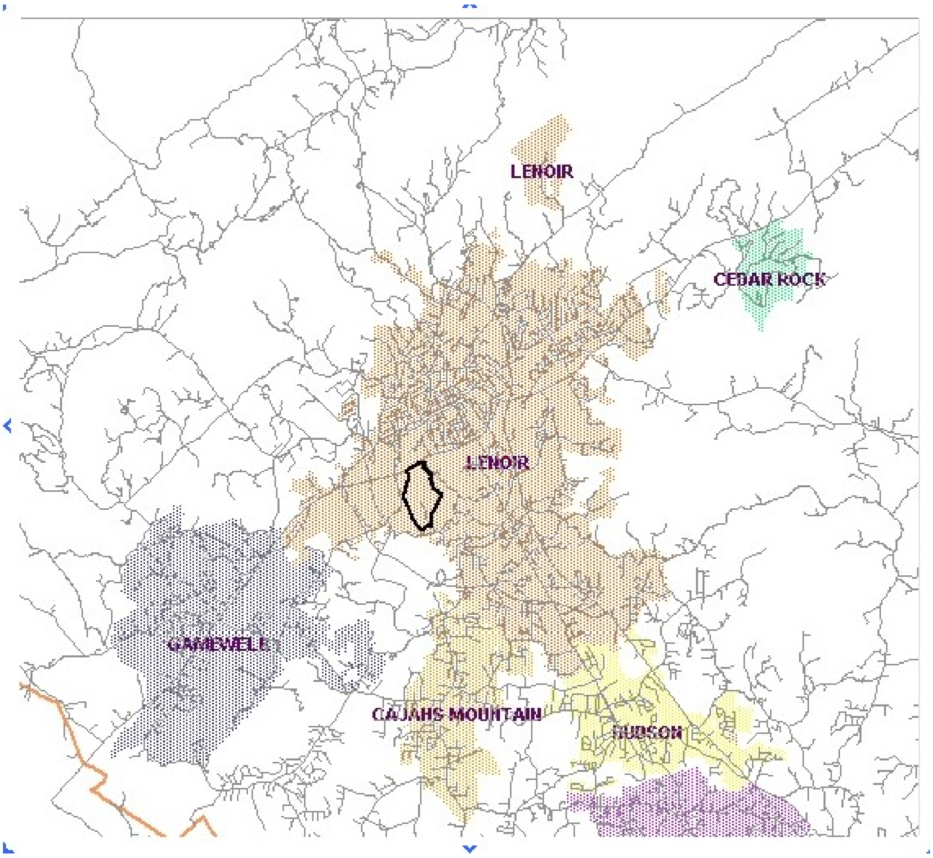

Figure 3:The city of Lenoir, NC, and the land Google acquired for the Lenoir data center. Source: http://maps.co.caldwell.nc.us/.

These data all support the image of Caldwell County's population being fairly typical for a rural American county with little ethnic diversity, higher-than-average levels of poverty, and a fairly high share of elderly residents. Economically, Lenoir has long been reliant on the furniture manufacturing industry that can be found in many areas in and around North Carolina. While some furniture manufacturing still takes place, many plants have closed in the past decades, with manufacturers opting to move to China, taking advantage of the low wages and economic boom there. This caused large-scale unemployment and economic collapse in some areas of North Carolina, similar to the problems present in the American Midwest, following the decline of the automotive industry there.

These problems how themselves in regional and state statistics on employment and the economy. The Western North Carolina Economic Index shows that, while there has been job growth in most of the 25 counties that make up Western North Carolina (WNC), the region remains behind the state average. The WNC area experienced, between 2002 and late 2006, a job growth of 5,75%, versus 8,2% statewide. However, the unemployment rate in the WNC area is the same as in the entire state of North Carolina, 4,8% (AdvantageWest 2006). This discrepancy of equal unemployment and unequal growth indicates that most of the jobs are created in other parts of North Carolina, most likely in the Triangle area (which consists of Raleigh, Durham, and Chapel Hill). This is in line with the rise of the Triangle as an area for high-tech industry, as well as the decline of the furniture industry in the state.

The decline of the furniture industry in the U.S. has been documented by Quesada and Gazo (2006), using statistics from several official and unofficial sources to determine the numbers of plant closures, as well as the reasons behind them. In their paper, they document the decline of both the primary and the secondary wood-processing industries. The first directly uses wood to create planks, beams, floors, etc. while the second uses those parts to create furniture and some other items. Thus, the two are closely linked, so that a decline in the furniture manufacturing industry immediately affects the primary wood-processing plants. Quesada and Gazo argue that the admission of China to the World Trade Organization (WTO) in 2000 greatly increased the competition that American furniture manufacturers faced from South-East Asia, resulting almost immediately in more plant closures and mass layoffs.

In total, Quesada and Gazo found reports of 168 plant closures, between 2000 and 2003, in non-official sources, across the entire United States. Official sources showed only 109 closures. The difference can be explained by smaller plants not being counted in official sources, and multiple plants on the same site being counted as one closure event. Nevertheless, it is safe to say that these closures dramatically affected North Carolina's labor market. In terms of plant closures, 73 out of the 168 non-officially reported plant closures took place there. For the United States as a whole, the time period 1999-2004 showed a 15% decrease in employment in the manufacturing industry

(Quesada and Gazo 2006). From these figures, it is easy to see that North Carolina is very strongly, and disproportionally, affected by the decline in furniture manufacturing.

This decline in furniture manufacturing still affects North Carolina. In 2006, the percentage of unemployed people in Caldwell County was 7,5%, much higher than the Western North Carolina average of 5,11%. While in 1995, Caldwell County had a lower unemployment figure than the WNC average (3,8% versus 5,95%), this advantage slowly eroded, until Caldwell was pushed over the WNC figure in 2001, shortly after China's admission to the WTO accelerated the decline of North Carolina's furniture industry (LINC 2007).

Google's new data center, then, will provide much-needed employment in the Lenoir area. The project, comprising of two data center facilities on a 215-acre site (see figure 4), will employ somewhere between 200 and 250 people, at an average salary of $48,000 (New York Times 2007). The jobs at Google's site will consist mostly of facility managers, hardware and software technicians, and system administrators. Google is currently helping local colleges to set up courses necessary for a job at their Lenoir site (Google 2008). However, the jobs to be created by Google in Lenoir do come at a cost. To ensure Google's choice for the Lenoir area, local and state authorities approved an incentives package amounting to more than $210 million worth of tax breaks and infrastructure improvements, spread out over the next 30 years (Businessweek 2007). The incentives package spawned opposition in North Carolina, for example by the North Carolina Institute for Constitutional Law, who are worried about increasing cost of the incentive packages needed to attract industries. A study conducted by the NCICL found that the Google incentives package accounts for over 50% of all incentives approved in North Carolina in the period 2004-2006 (NCICL 2007).

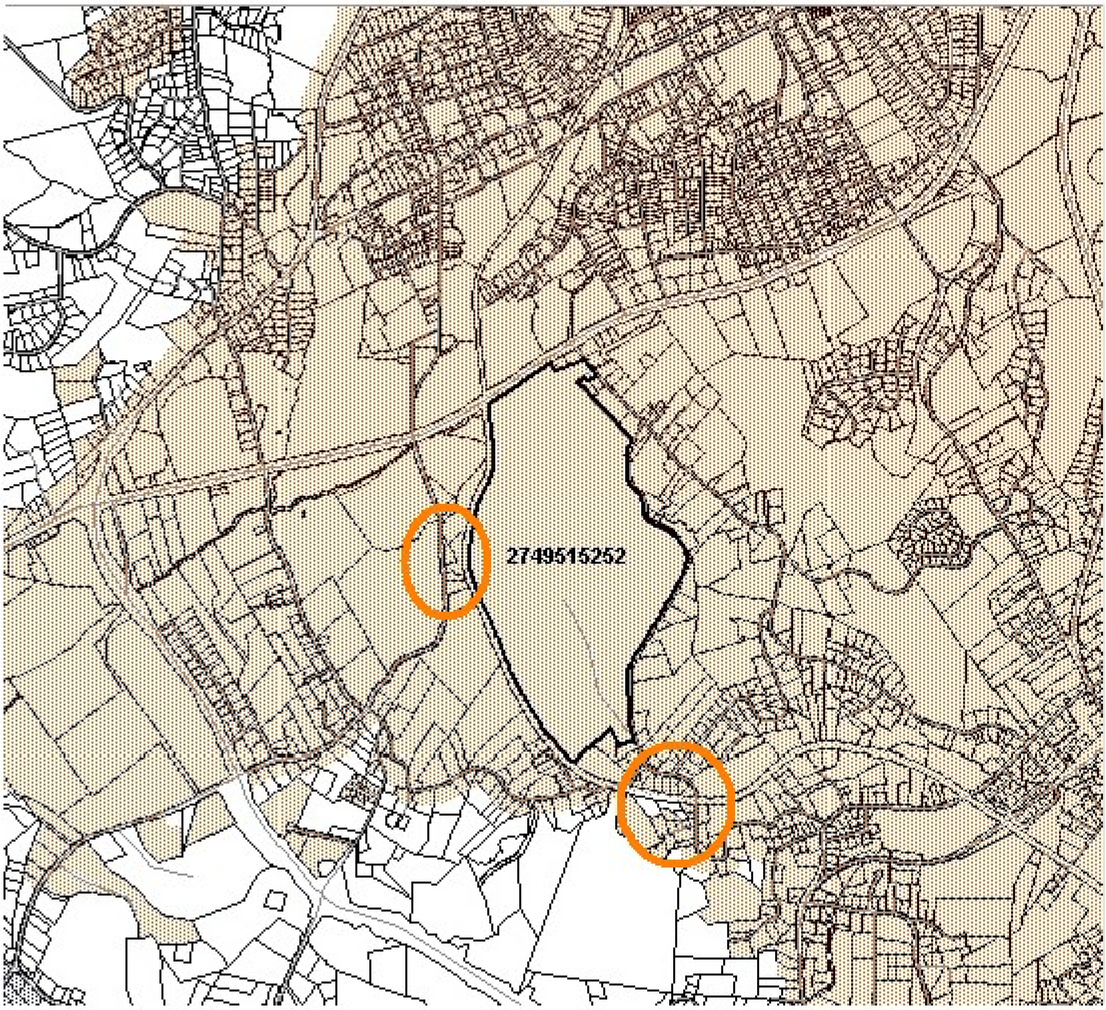

Figure 4: Google's 215-acre site in Lenoir, NC. The circled areas contain several much smaller plots of land also owned by Google. Source: http://maps.co.caldwell.nc.us/

As explained in chapter 4, fiber optic connections are one of the two most critical resources of a data center. Bandwidth on these fiber optic cables is often thought to be lacking in rural areas. This is supported by a Pew Research report, which shows that consumer-level broadband has only been adopted by 31% of all residents of rural areas in the United States, compared to (respectively) 52% and 49% for urban and suburban areas (Pew Internet 2007). However, a document from regional economic development group AdvantageWest (2008) shows that all counties in the Western North Carolina (WNC) area have some form of broadband access available through multiple service providers. Caldwell County has eight providers offering internet, telephone and/or television services, out of thirteen total providers operating in the WNC area (Advantagewest 2008). As of December 31st, 2006, slightly less than 70% of all households in Caldwell County had the possibility of subscribing to a high-speed internet connection (e-NC 2007).

Besides internet access on the consumer level, there are also documents showing the presence of local and regional fiber optic networks in the WNC area. While regional data on the presence of commercial networks is not publicly available, there are several non-commercial networks that indicate a certain level of internet connectivity. These networks are operated by BalsamWest, e-Polk and ERC Broadband. The presence of these networks indicate that the WNC area is well-connected to the wider fiber optic infrastructure (see figure 5). This is backed up by a statement from Lenoir planning director Chuck Beatty, who stated in an email that, as far as he knew, the presence and extent of the fiber optic network in the area did not play a role in Google's location decision (Beatty 2008). This could mean that Google deemed the local networks more than adequate for its purposes, but another possibility is that Google will create its own fiber optic network, either by laying down its own cables, or by buying up unused capacity of existing lines[7].

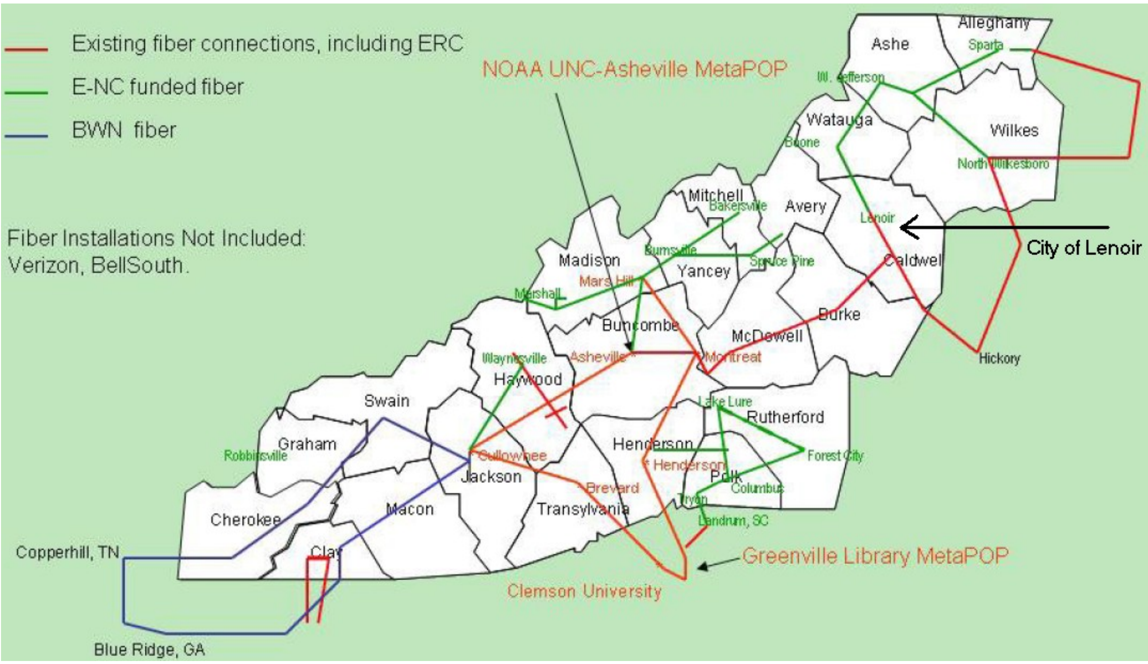

Figure 5. Non-commercial fiber optic networks in the Western North Carolina (WNC) area. Source: AdvantageWest. Edited by author.

One important reason that explains Google's interest in the Lenoir area is the decline of the furniture manufacturing industry itself. Because of the size of the industry, it placed considerable demand on the regional infrastructure. With furniture manufacturing in decline for quite some time now, resources in the region were underutilized. This was confirmed as being a factor by Chuck Beatty (2008), Lenoir's planning director. Because of the lower power consumption in recent years,

[7] There were rumors of Google buying this so-called 'dark fiber' around 2005-2006. However, there is no concrete evidence of Google creating its own fiber optic network.

after furniture manufacturers moved away or closed down, Google could step in and take advantage of the low cost of utilities. Considering the large amount of electricity needed to power the facility (see chapter 4), this is likely to have been a major factor in Google's location decision. In fact, it was Duke Energy, a major North Carolina electricity provider, who helped Google find the Lenoir location. Duke Energy is the operator of the Marshall coal power plant in Catawba County, as well as the McGuire nuclear power plant in Mecklenburg county. These plants are the two major power plants that are closest to Caldwell County/Lenoir, and the WNC area in general. Thus, these plants are likely to have seen the biggest drop in electricity consumption following the decline of the furniture industry. Since it is inefficient for a power plant to continually operate at less than maximum levels (especially for nuclear plants), it is not surprising that Duke Energy would help attract a new, major customer to the area. An example of this is Duke's monitoring of Google's negotiations with state and local governments. They also sold 60 acres of land to Caldwell County, whose officials were assembling the 215 acre site where the data center is now under construction (News & Observer 2007).

Concluding this chapter, it appears that, while there are several factors at work that explain Google's location decision, they are all connected to the decline of the furniture industry, documented by Quesada & Gazo (2006). Because many furniture manufacturing plants in the Western North Carolina area either closed down or moved overseas, consumption of electricity dropped, allowing Google to step in and take advantage of the surplus capacity, assisted by Duke Energy, a major player in the North Carolina electricity market. Cheap land is another factor, and while the low cost is partly due to Lenoir's rural location, the decline of the furniture industry is a factor here as well. These two factors were confirmed to be part of Google's reasoning by Lenoir planning director Chuck Beatty (2008), as opposed to fiber optic capacity in the region, which is not a factor, most likely because the regional networks are already adequate for Google's purpose. Furthermore, the unemployment in the region, combined with Google setting up programs at regional colleges, ensures that Google will have access to the needed personnel to run the facility. The final factor in Google's decision is the $210 million incentives package provided to Google by state and local authorities (NCICL 2007). However, if, and to what extent, this package was critical to Google's decision remains unknown.

6. Conclusion

In this paper, location theory for data centers has been explored, against a background of Google's new Lenoir, NC data center, as well as a comparison with location theory for general IT businesses. In the Netherlands, these general IT businesses seem to follow patterns based on customer location and general accessibility, either on a national or a regional scale (Atzema 2001). The United States IT sector is located mostly in certain metropolitan areas, with the highest concentrations in and around Boston and San Jose. The location factors here seem to be mostly historical, and the IT businesses in the United States appear to have evolved out of older high-tech sectors (Cortright & Mayer 2001). This makes it difficult to point out any specific location factors, although there still is a tendency to locate near population and economic centers (near 'the market').

Data centers are somewhat different from these general IT industries. They have needs that manifest themselves in the internet's infrastructure, focusing on virtual access instead of physical access, meaning that data centers prefer to be accessible through the internet. Still, the need to locate centrally within the larger physical area exists, and data centers still often locate near major cities. Usually, however, this is because that is where virtual accessibility is also greatest: major cities are better connected to the internet compared to rural areas. However, Malecki (2002), as well as more recent backbone maps like those from network providers Cogent (2008) and Level3 (2006), show that not all cities in the United States are equally connected to the internet, indicating that some cities make better locations for data centers than others. Combined with the information on datacentermap.com (2008), this shows that data centers tend to locate near fiber optic hubs like Denver and Miami. However, when it comes to Google's Lenoir site, the rural location of the area springs to mind immediately, in sharp contrast with the more common urban data center locations described by Evans-Cowley (2002). When it comes to the specific Lenoir location, however, the fiber optic connectivity of the area was not an issue. Instead, the most important location factors were the resources available in the area, like cheap and available land, water and electricity (according to Lenoir planning director Chuck Beatty (2008)). Since the decline of the furniture industry in the area (Quesada and Gazo 2006), many resources were underutilized, and made it possible for Google to step in to take advantage of the cheap electricity and the other resources in the area their data center needs (Beatty 2008). Unfortunately, requirements that are specific to Google's data centers could not be identified. Partly this is due to Google's secrecy regarding its operations, and partly this may be because Google has little requirements for its data centers beyond the general ones like electricity costs. However, as discussed in chapter 4, Google has spread its data centers across the world and the United States, and any new data center locations, like Lenoir, may be chosen to create a network that optimizes the accessibility of Google's services.

On the topic of government assistance, the state and local governments in North Carolina were aided by regional electricity provider Duke Energy, who acted as a monitor to the talks with Google, and helped find the Lenoir location (News & Observer 2007). Combined with the $210 million tax incentives package offered to Google by state and local governments (NCICS 2007), and efforts by the city of Lenoir to piece together the plot of land that Google ultimately bought to construct its data center (News & Observer 2007), it is clear that the various governments of the area were active participants in the process of setting Google up in Lenoir. Thus, the Lenoir location decision was mostly based on the availability and price of the needed resources to build and operate the data center, as well as the presence of cooperative local and regional governments. The availability and quality of any regional fiber optic networks does not appear to have played a significant role in Google's location decision.

Sources[8]:

Abbate, J. (1999), *Inventing the Internet*. Cambridge, MA, Massachussetts Institute of Technology.

AdvantageWest (2006), *Western North Carolina Economic Index*. Accessed on 15-12-2007 through http://www.advantagewest.com/uploads/File/WNC_Economic_Index_August2006.pdf

AdvantageWest (2008), *WNC Regional Telecom Broadband Service Providers*. Accessed on 23-03-2008 through http://www.advantagewest.com/uploads/documents/WNC%20Regional%20Telecom%20Broadband%20Service%20Providers.pdf

Atzema, O. (2001), *Location and Local Networks of ICT Firms in the Netherlands.* In: Tijdschrift voor Economische en Sociale Geografie vol. 92 no. 3 (2001) pp. 369-378.

Baseline Magazine, (2006), *How Google Works*. Accessed on 24-11-2007 though http://www.baselinemag.com/article2/0,1540,1985040,00.asp

Beatty, C., (2008), emails exchanged with author.

Businessweek (2007), *The High Cost of Wooing Google*. Accessed on 26-03-2008 through http://www.businessweek.com/magazine/content/07_30/b4043066.htm

Camagni, R. and R. Capello (2005), *ICT's and Territorial Competitiveness in the Era of the Internet.* In: Annals of Regional Science vol. 39 no. 3 (2005) pp. 421-438.

Cogent Communications (2008), *Network Map*. Accessed on 23-01-2008 through http://www.cogentco.com/htdocs/map.php

ComScore (2008), *ComScore Releases 2007 U.S. Internet Year in Review*. Accessed on 20-02-2007 through http://www.comscore.com/press/release.asp?press=2043

Cortright, J. and H. Mayer (2001), *High Tech Specialization: A Comparison of High Technology Centers*. Portland: Center on Urban & Metropolitan Policy.

Datacenterknowledge.com (2007), *Google Data Centers: $3,000 A Square Foot?* Accessed on 07-01-2008 through http://www.datacenterknowledge.com/archives/2007/Nov/12/google_data_centers_3000_a_square_foot.html

Datacenterknowledge.com (2008), *Google Data Center FAQ*. Accessed on 01-04-2008 through http://www.datacenterknowledge.com/archives/2008/Mar/27/google_data_center_faq.html

Department of Energy (2006), *U.S. Average Monthly Bill by Sector, Census, Division, and State 2006*. Accessed on 07-01-2008 through http://www.eia.doe.gov/cneaf/electricity/esr/table5.xls

Department of Energy (2007), *Average Retail Price of Electricity to Ultimate Customers by End-Use Sector*. Accessed on 23-01-2008 through http://www.eia.doe.gov/cneaf/electricity/epa/epat7p4.html

e-NC (2007), *2006 High Speed Internet Access*. Accessed on 27-03-2007 through http://www.e-nc.org/pdf/2006_HSIA.pdf

[8] Dates in DD-MM-YYYY format.

Evans-Cowley, J. (2002), *Telecom Hotels – A Planner's Guide*. American Planning Association Report 505.
&
Evans-Cowley, J., E.J. Malecki and A. McIntee (2002), *Planning Responses to Telecom Hotels: What Accounts for Increased Regulation of Co-Location Facilities?* Journal of Urban Technology vol. 9, no. 3, pp. 1-18.

Google Inc. (2007), *Google's Goal: Renewable Energy Cheaper than Coal.* Press Release. Accessed on 29-11-2007 through http://www.google.com/intl/en/press/pressrel/20071127_green.html

Google Inc. (2008), *Lenoir Data Center*. Accessed on 26-03-2008 through http://www.google.com/datacenter/lenoir/index.html

Halper, R.F. (1985), *Computer Data Center Design*. New York: Wiley-Interscience.

Hancock & Algozzine (2006), *Doing Case Study Research*. New York: Teachers College Press.

Hafner, K. and M. Lyon (1996), *Where Wizards Stay Up Late: The Origins of the Internet.* New York: Simon & Schuster.

Harris, S. and E. Gerich (1996), *Retiring the NSFNET Backbone Service: Chronicling the End of an Era*. Last accessed through http://merit.edu/networkresearch/projecthistory/nsfnet/nsfnet_article.php on 06/02/2008.

Hillstrom, K. (2005), *Defining Moments: The Internet Revolution*. Detroit: Omnigraphics.

Level3 Communications (2006), *The Level3 Network*. Accessed on 23/01/2008 through http://www.level3.com/images/global_map/Level_3_Network_map.pdf

Koomey, J. (2007a), *Estimating Total Power Consumption by Servers in the U.S. and the World*. Accessed on 07-01-2008 through http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf

Koomey, J. (2007b), *Data Center Electricity Use: What We Know.* Accessed on 02-12-2007 through http://www.energystar.gov/ia/products/downloads/koomeysummaryondatacenterpower.pdf

Maggioni, M.A., (2002), *Clustering Dynamics and the Location of High-Tech Firms*. New York: Physica-Verlag.

Malecki, E.J. (2002), *The Economic Geography of the Internet's Infrastructure*. In: Economic Geography, vol. 78 no. 4 (2002), pp. 399-424.

Milligan, L. (2001), *Local Governments and the High-Tech Economy*. In: Special Data Issue no. 3 (2001). Washington DC: International City/County Management Association.

Moss, M.L. and A.M. Townsend (1998), *Spatial Analysis of the Internet in U.S. Cities and States.*

LINC database (2007), *Employment & Income*. Accessed on 16-12-2007 through http://linc.state.nc.us/

North Carolina Institute for Constitutional Law (NCICS) (2007), *The Incentives Game: North Carolina Economic Development Incentives*

News & Oberver, 16-02-2007, *Duke Energy Helped Land Google*.

New York Times, 15-03-2007, *Google Is Reviving Hopes For Ex-Furniture Makers*.

Pew Internet and American Life Project (2007), *Home Broadband Adoption 2007*. Accessed on 23-03-2008 through http://www.pewinternet.org/pdfs/PIP_Rural_Broadband.pdf

Quesada, H.J., and R. Gazo (2006), *Mass Layoffs and Plant Closures in the U.S. Wood Products and Furniture Manufacturing Industries*. In: Forest Products Journal vol. 56, nr. 10 (2006).

Schaeffer, H. (1981), *Data Center Operations* (1st ed.). Englewood Cliffs, NJ: Prentice-Hall.

Search Engine Optimization Project (2008), Google PR. *About Google datacenters*. Accessed on 11-02-2008 through http://www.seop.ru/en/datacenters-2.html

Sommers, P. and D. Carson (2000), *Ten Steps to a High Tech Future: the New Economy in Metropolitan Seattle*. Accessed on 02/12/2007 through http://www.brookings.edu/dybdocroot/es/urban/sommers/sommersreport.pdf

Threadwatch.org (2006), *Google Has 41 Data Centers After All*. Accessed on 24-11-2007 through http://www.threadwatch.org/node/7463

Vise, D. and M. Malseed (2005), *The Google Story*. New York: Bantam Dell.

Washington Post (2003), *Dissertation could be a Security Threat*. Accessed on 19-01-2008 through http://www.washingtonpost.com/ac2/wp-dyn/A23689-2003Jul7