Research at the Edge of the Network

Greg Kochanski, Chilin Shih
{gpk,cls}@bell-labs.com

December 4, 2001

Abstract

We make the case that Lucent will likely need some research relating to things at the edge of the network. Later sections describe research time scales, the appropriate scope of the research area, and evaluation metrics.

1 The edge of the network

The edge of the network is currently less developed than the middle. With modern optical networks, we can transport almost more bits than we can imagine. The bottleneck to the network is the edge.

The bottleneck is the edge both on the level of network hardware (most people still connect via 56k modems) and at the level of applications. Other than watching TV, no one knows what people would actually *do* with multi-megabit connections. Wireless is part of the bottleneck, and it is certainly a good field for Lucent, but even there, no one knows what people would do with a megabit connection to a tiny screen.

We need research related to the edge of the network, because the edge of the network controls the demand for bits. Even if we do not build the edge of the network, we need to know what is coming.

1.1 Wireless terminals

Lucent is limited in its ability to define wireless offerings, because we do not control both ends of the wireless link.

Since we have little influence on the handset side, we are pretty much limited to supporting industry standards: we have no chance of forcing people to accept our standards. This is a serious limitation. By being restricted to following the standards we have several problems:

- Our equipment supports the same features as our competitor's. Therefore, we are selling something that is close to an interchangeable commodity.

  Commodities have low profit margins, because if one can't provide better features, one's best hope of attracting customers becomes offering a lower price than the competition.

- Our products start their development cycle behind the competition. The companies that are driving the standards have advance warning on what the standards might be. Thus, they may be able to ship products earlier than Lucent.

  First-to-market tends to get a larger market share and (at least initially) higher profit margins. If we are late, we get lower profit margins.

- If the standard has flaws, we may be building a system that will not work well enough and/or may not be purchased.

  For example, high-quality voice over IP may not be easily achievable for the current standard. If so, being locked in could be a serious problem.

Lucent needs to have some influence on handsets so we can have more influence on the standard development. Either we build them (unlikely, given our history) or we partner closely with handset builders.

Either way, we need some research relating to handsets (and the people who use them). We need to have a patent portfolio to help us control the relationship with our partner, and we need to understand what is happening in the technology.

1.2 Talking to machines

In the late 1990s, speech synthesis and recognition technologies reached the point where a computer could speak well enough to be understood by the general population, and computers can recognize task-specific speech often enough so that the mistakes can be repaired in a well-designed dialogue. As a result, we are seeing a rapid deployment of dialog systems. There are now dozens of systems publicly operated by major corporations, doing complex functions.

Potentially, these dialog systems are crucial to pocket-sized devices, where the user does not have much of a screen. The success of Palm pilots and related hand-held devices has shown how much people value portability and compactness. We think it is not prudent to focus Lucent's research exclusively to support people who will carry laptops and use data-only services.

While the prospect of automated "personal assistants" was much hyped in the last decade, and is not yet a near-term prospect, such services will become feasible in the foreseeable future. They might well become an always-on application that alters the way people interact with each other and the network.

Dialog systems are a growing technology, one that Lucent needs to track, but also a technology with a set of serious unsolved problems. Generally speaking, the field has been pushed into engineering practice and commercial systems while some of the fundamental science is not yet understood. Current systems are limited because of this lack of understanding. There are two related fields: linguistics and speech modeling and natural language processing.

For example, we do not know how people control dialogs. We don't know details of what information is carried by prosody, or how much. We don't know how to respond if the user says anything unexpected.

We need to expand our IP in the area, making sure we have a reasonable patent portfolio so that we can compete later. We also need to take the opportunity to fill some of the gaps in our understanding that have been uncovered as we have built systems. Basic research in this area, especially modeling of neuromuscular control, could plug directly into the synthesizer and recognition parts of a dialog system, and result in easier, clearer user interactions. Better understanding of what is being communicated by prosody could lead to much more natural, less machine-like interactions.

1.3 Lucent also supports equipment

Today's conventional model of Lucent's business is roughly this: Lucent makes the boxes, and our customers (the service providers) deal with end-users. In that model, one would think that Lucent has little use for user-interface research.

However, Lucent needs to service one very important class of users: the people who install, provision and repair our products. The research portfolio should include fields that are necessary for this purpose.

Lucent's revenue from installation, service, upgrades and maintenance is a respectable fraction of our total revenue. Anything we can do to reduce the cost is valuable.

A few possibilities come to mind with their associated technologies:

- Database mining in service records to find defects and improve products (Natural language processing).
- Expert systems to aid service personnel in finding problems (Machine learning techniques, Natural Language Processing, techniques from mixed-initiative dialog systems).
- Machine generated indices to product manuals in foreign languages, to allow non-English-fluent technicians more access to the documentation (Machine translation, automatic summarization).
- Tools to help service providers deal with customers. One example might be a system to measure the health of wireless handsets, store the data, and report it to service providers. AT&T, for example, apparently doesn't know if the customer's handset is getting worse service than average. (statistical modeling, physics, expert systems).

1.4 Image processing and video coding

Image processing and video coding will soon be relevant for wire-line applications. Certainly, many people will download and watch videos, even if society never accepts video phones. They will watch them on a variety of different-sized screens, over links with different characteristics, and they will demand high-quality video. Eventually, it will be relevant for wireless links. Video coding technology is relevant because it is the only in-hand technology where its broad use could dramatically increase the number of bits flowing through the network.

2 The scope of research

A reasonable definition of "research" is "the people that solve problems which might be required for future generations of products."

For a research project to be useful, the problems it solves should:

(a) have relatively broad scope so that they have a larger chance of being used in a future product,

(b) should develop expertise in techniques that have a broad scope of application to potential products,

(c) develop patents directly relevant to Lucent or which might be exchanged for other company's patents that Lucent needs, or

(d) other less-direct justifications such as public relations (both to customers and to potential employees), opening and advancing new fields potentially relevant to Lucent, and helping Lucent to understand some technological opportunity.

Fundamentally, research management is a game: There are research projects you can throw at a moving target. One needs to throw towards where the target might be when the project finishes. It always helps if the projects cover a broad area, but an accurate throw only helps if you know where the target will be.

2.1 How far ahead must research look?

The the research time scale has relevance to all four points.

The time scale from initiation of a research project to initial product announcement is unavoidably long. Lambda Router, for instance, is touted as having an 18 month time scale, but that time just dates from the formal beginning of the project. Research here related to the Lambda Router dates back as far as the hiring of Suzanne Arney (c. 1993) which began Lucent's MEMs program, and also has roots in waveguide multiplexers (1996 or before) in Alice White's department, and work on free space optical interconnects between chips which was done before 1994.

It would have been impossible to build Lambda Router as fast as it was done without some of the prior work. From a clean slate, it could not have been done in less than three years. More importantly, no one would have had the confidence to push the project forward if Lucent had been starting from a clean slate. Lucent might well have waited until it was proven possible by some other company.

Even in software areas, where it might be assumed that research can transfer more directly into products, there are not really any shortcuts. Development organizations invariably end up rewriting research code. We've been told by two separate sets of development organizations that "We don't want your code, we want your algorithms." These statements were not results of turf battles, they were honest statements from developers who were trying to get products out the door as fast as possible. This statement doesn't reflect poor research, either: its basis is that research code and product code fill different needs.

Research code (code written as the problem is being defined) has to be flexible, easy to change, and it has to be simple enough for the coder to understand the effect of changes. On the other hand, product code has to be reliable, and cannot be allowed to fail in a way that costs the customer money (*e.g.,* crash or corrupt data). As a consequence, it has extensive error checking (50% of source lines are related to error checking and recovery in the LSS product speech synthesizers), it is written formally with specialized data types, which can too much time to generate in a research environment. Product code may also need to be written with more of an eye for efficiency, and may need to be multi-threaded or have other constraints that would make it too hard and slow for the researchers to to succeed in their main task: finding and testing the right algorithm.

I will assume a three year time scale for the time from the initiation of research to product release. This is a compressed, lower limit: one year for exploratory work, one for applied research, and one for actually building the product.

Pushing the product cycle to be too short has its dangers, too. It leaves fewer opportunities to think and test, and more opportunities for mistakes to pass on from research to development, and development to the final product. Mistakes are always more expensive to catch and fix in the later stages of a project.

2.2 Stability of corporate goals

Historically, the corporation's goals have not been stable over a three-year window.

- Three years ago, Lucent was an optical networking company; today Lucent is primarily a wireless company.
- Three years ago, our target customers were dot-coms; today they are big telephone companies.
- Six years ago, data networking was unimportant except for faxes; today it is perceived as dominant.
- Six years ago, we made products that cost less than \$100, and sold them to millions of people; today, we make products that cost more than \$1,000,000 and sell them to fewer than 100 customers.
- Nine years ago, we were a service provider; today we sell equipment.

Therefore, any research plan has to anticipate that the company will change its goals (to a substantial degree) before the research comes to fruition. This variability in corporate goals makes it impractical to focus research too tightly.

Tightly focussed research will often be rendered irrelevant by changes in the corporation. This sets a lower limit on the breadth of the scope in points **(a)** and **(b)**, above.

2.3 Patents as enablers: point *c*.

Lucent has historically attached modest importance to patents. In previous evaluations of research, the value of the intellectual property that was generated was set strictly by the licensing revenue it brought in. This would be the correct answer if there were an efficient market in patents, but there isn't.

There is nothing to price patents on the open market, like a stock exchange. The rights to use most patents are traded in large blocks, in exchange for rights to large blocks of other patents. Only a few are valued in terms of cash, and the negotiations are often slow and complicated.

As a consequence, the licensing revenue from Lucent's intellectual property is only a part of the total value to the company. The patents also supply value by allowing the company to legally produce its products. Many of Lucent's current products would be in violation of some patent or the other, if we had not swapped patent rights. Patents then are enablers of much of Lucent's revenue and profit.

One way to put a dollar figure on this "enabling" is to imagine Lucent without a patent portfolio and then consider how much money we would have to pay to license all the patents we use. The cost would clearly be very large.

This cost applies technology by technology: if we had no patents in some technology (*e.g.,* Ge-doped optical fibers), and we needed to incorporate this technology into the company's business plan, we would have to pay cash to license any patent we would use, because Corning (the company that holds some crucial patents) would be unwilling to swap. In reality, Corning allowed Lucent to use their patents, even though they were our direct competitor, because Lucent had a strong enough portfolio in that field to be valuable to Corning.

Likewise, Lucent cannot expect to enter a new field and rapidly develop its patent portfolio. Other than narrow patents that protect specific products, many of the relevant patents will have already been filed by others, years before. It takes years for a technology to develop from where the possibilities are clear enough to file broad and important patents to where one can actually build products and make money.

As an example, consider two Bell Labs patents probably infringed by Microsoft Windows (4,700,320 and 4,710,761). The patents were applied for in 1985, and it wasn't until c. 1992 that there was a large market for computers with graphical user interfaces. Similarly, Public Key Cryptography (the RSA algorithm, US patent 4,405,829) was filed in 1983, but this technology only really became commercially important in c. 1997 when people started buying things over the Internet.

Thus, there is considerable value to Lucent for Research to develop and maintain a patent portfolio in any technology that is likely to be relevant to Lucent in the next decade. The value comes from avoiding the need to license patents that Lucent needs to operate.

2.4 Predicting the future

People are well known to be bad at predicting the future impact of technologies. This is part of the broader problem of predicting the future, and modern techniques like spreadsheets and business plans are only modest advances beyond chicken entrails and tea leaves.

A historical view shows an amazing variety of incorrect and missed predictions from reputable experts. For instance, in the late 1940's, people seriously believed that future commuters would get to work in personal airplanes and helicopters. On the other hand, no one managed to predict the extent that the automobile would change American society. Closer to Lucent, no one predicted the vast number of home pages that sprang up on the Web until they happened.

More relevantly, even in one company, it has proven impossible to predict which products will be important. For example:

- Wireless research was shut down by AT&T, because the market for cell phones was predicted to be small.
- C++ development was shut down three times, losing us the opportunity to sell the first compiler. C++ is still the industry standard computer language, but AT&T never made a dime on it.
- High speed networking development was shut down because management believed that no one would need it. Lucent has recently suffered severely from that decision.
- 2.5 Gb/s optoelectronic research was briefly shut down back in c. 1990, again because no one expected to actually need anything more than 1Gb/s in a single fiber.

These are not examples of below-par management: they just show that no one can predict the success or failure of technologies more than a year in advance.

Unfortunately, this inability to predict means that if one focuses research, it is likely that it will be focussed on the wrong goals. This is yet another reason that the scope of research might as well be fairly broad.

3 Which comes first, the applications or the technology?

The value of research frequently shows up in unexpected places as new applications materialize. This happens when a problem is solved in a fundamental way, and can be applied over a wide domain. One recent examples is the application of optical character recognition (OCR) to network security. Henry Baird is a leading expert in OCR at Bell Labs who left for Xerox Parc soon after the trivestitute. (OCR research was relevant to a service provider like AT&T, though its value was less obvious for an equipment company like Lucent.)

As the Web developed, companies began offering free access to large databases or free disk space. This was intended for human use, often as part of an advertising package. However, people started making automated use of these services. Programs would register for hundreds of free e-mail accounts at Yahoo and Hotmail in order to get access to the free disk space and anonymous e-mail accounts. Automated scripts can copy web databases, item by item, essentially stealing the investment of the database creator. It then became necessary to offer services only to humans. Baird understood enough about the human visual system and OCR systems to realize that he could construct a reverse Turing test. An elegant adaption of techniques he developed for training OCR systems allows a server to produce queries that are trivially easy for humans to answer, but impossible for automated systems to decode. The system is now commercially used on Yahoo.

Such an adaptation takes little effort when the researcher has the right knowledge and sees the right need. Usually, we justify research is justified by pointing to a well known application, and aiming the research at the target. But, this is an example where the research came first, and there was no obvious way of justifying the research inside Lucent until the application appeared. This example shows that if one neglects curiosity-driven research and forces all projects to be justified by applications that the researchers and management can imagine, Lucent will lose some technologies, when the technology arrives before the application.

4 The scope of lucent

Focusing Lucent into a narrow range of products may be a necessary thing for short-term survival, but in the longer term it puts the company in a position where its revenue will be extremely volatile. Two separate effects increase volatility: First, the narrow focus means that the company's revenue rises and falls with the spending of single homogeneous set of customers. Second, supplying hardware with a long lifetime is an intrinsically variable game: if they ever have too much capacity, your customers can simply stop ordering and use the equipment they have.

As an example, suppose that it was believed by the service providers (our customers) that demand for telecommunication services would drop by 10% over the next year. They would then need to order very little from us this year. If the network ever becomes over-provisioned in reality, we can expect severe drops in revenue. Even in more likely, less drastic cases, one expects that our revenue will be more variable than that of the service providers.

An analogous company to Lucent is The Boeing Company. Boeing, after 1970, built commercial aircraft as its primary revenue source. Commercial aircraft are, like Lucent's products, technologically sophisticated, expensive, ordered by a few large customers, and long-lasting. In the 1970 recession, Boeing's orders dropped dramatically because airlines had all the airplanes they needed to carry a reduced passenger load.

> Unfortunately, due to the recession in the aviation industry, Boeing went 18 months without a single new domestic order. ... In the Seattle area alone, the Boeing workforce was cut from 80,400 to 37,200 between early 1970 and October 1971. Thousands of former Boeing employees, finding little in the local job market, looked for work elsewhere. Things became so bad in Seattle that a billboard on the city's outskirts read, "Will the last person leaving Seattle turn out the lights"

From official Boeing web pages, http://www.boeing.com/companyoffices/history/boeing/markets.html

The same thing could happen to Lucent. Indeed, it could be argued that Lucent's current troubles are a similar example. Service provider orders have dropped 15%, instead of an expected growth of 10% or more, and Lucent has faced some painful adjustments. In turn, Lucent has cut its semiconductor orders to Agere even more severely (because Lucent has too much unsold product). Thus, one step further down the food chain, Agere is suffering more severe decreases of revenue, and larger, more damaging layoffs.

To reduce this kind of damaging instability in the future, Lucent needs to expand its scope and diversify. Going forward, we need more than just one homogeneous set of customers who will all cut orders at the same time. Clearly, it was a mistake in 1998-2000 to diversify into a new customer base, but it was only a mistake because we chose a customer base that was unstable and, in the end, didnt́ actually have any money.

> Condit is moving Boeing's headquarters from Seattle to Chicago, Dallas or Denver later this year, partly to reflect that it's no longer simply an aircraft builder. ... Since last year it's started the Internet service, a financing arm and a satellite-based air-traffic management unit.

Apr 24, 2001 Airwise News, http://news.airwise.com/stories/2001/04/988137769.html

Boeing, for example, has diversified since its brush with disaster and expects to ride out the current slump. in commercial airliner orders with a smaller fraction of layoffs than Lucent has announced. Boeing currently gets revenue from three main units - military aircraft and missiles, commercial airplanes, and space and communications; each from a different set of customers.

For a simple measurement of diversification, I selected 16 large industrial companies: the members of the Fortune 100 that largely make capital equipment[1]. Diversification was measured by counting words[2] in the http://finance.yahoo.com corporate profile. After taking out the businesses that are being spun off as Agere (5 words), Lucent had 21 words, with only Ford (20 words) and Compaq (16 words) having a simpler description[3]. Both Ford and Compaq are both very focussed companies with few distinct products and few classes of customers. This numerical conclusion is supported by reading the text: Lucent indeed has fewer distinct lines of business, and fewer classes of customers than "comparable" companies.

Diversification carries its own risks, of course, but looking out three to ten years, the time relevant for research planning, I expect that Lucent will not be as tightly focussed as it is now, in late 2001. Thus, research may want to support more than the current Lucent.

[1]BA, CAT, CPQ, HON, HWP, IBM, INTC, F, GE, GM, LMT, LU, MMM, MOT, UTX

[2]ignoring "a", "the", "in", and similar helper words.

[3]The average length of description of companies being compared to Lucent was 28 words