

Lecture 5- Pre-lecture reading

Cellular economics: the “million-dollar cell”

The cost of a cellular system is obviously a complicated calculation, but we can simplify it dramatically for our strategic purposes. The overwhelming majority of the investment in any cellular system is in the cost of those “cell towers” (more accurately “base stations”) we must install in the center of every cell. We would normally expect the cost of a base station to depend on how many radios we have in the cell, but in fact the cost of the radios is a very small part of the total cost of a base station. Fixed costs include the land (including legal fees associated with planning, applications for zoning variances, and legal appeals when those applications are denied), a building to house the equipment, an expensive mast with antennas and thick coaxial cables, maintenance, control, and information systems, connection to the electric grid and back-up power and wireline telephone connections back to the switch. The result is an average cost of about \$1 million per base station. See slide 3, below.

The Million Dollar Cell

Mostly “fixed” costs (not just radios)



Legal fees for local approval*
Land and improvements
Engineering planning*
A building
Mast, antennas, cables
Radios, combiners
Control equipment
Maintenance equipment
Information systems
Connection to electrical power
Backup power
Connection to central switch

***These “people costs” are significant and are included in the cost of a cell**

Copyright 2009: Richard H. Frenkiel

3

Because all these other costs are so high, **the cost of a base station** is essentially independent of how many channels (radios) the base station uses. On the other hand, **the capacity of a base station** is entirely dependent on the number of radios (which obviously depends on the number of channels we get from the FCC). This leads to a strange economic structure that significantly favors a monopoly. If the FCC gave all the cellular channels (which turned out to be about 800 channels) to AT&T they would create an AT&T monopoly.

Some channels would be needed for signaling and control and varying user densities made cells less efficient, but the monopoly would serve about 2000 users per million-dollar cell, so it would cost AT&T about \$500 per user to build the system. As we will calculate next time using financial reports, this would result in users paying about \$30 per month for cellular service.


If the FCC wanted competition and divided the channels between two competing service providers (which they did), they would each get only 400 channels. Their systems would only serve about 1000 users per million-dollar cell and their system cost would therefore be \$1000 per user. As we will calculate in the next lecture, users would have to pay about \$50/month.

If the FCC wanted a lot of competition and divided the channels among four competing service providers, their systems would only serve about 500 users per million-dollar cell. The systems would therefore cost about \$2000 per user and we will calculate that the users would have to pay more than \$80/month.

AT&T argued that a monopoly was justified because it led to much lower cost per user, but AT&T's opponents would argue that competition would lead to the best results in the long term. AT&T's argument was mathematically correct, but as we will see in later lectures, the argument for competition was also correct. Competition tends to keep profits (and prices) down.

As we have mentioned (and will discuss in Lecture 7), the FCC didn't give AT&T a monopoly. Instead, they authorized two systems per location, each with about 400 channels. They probably would have preferred four systems to create more competition, so the fact that they limited competition to two systems suggests that they recognized the cost advantage that AT&T was calculating and struck a compromise. A decade later they would authorize additional channels at a slightly higher frequency (2 GHz) to create two additional competitors, but those new competitors would have to pay a significant price (about \$500M) for an amount of spectrum that had been free for the original systems. As we will see in lecture 8, this put them at a significant economic disadvantage relative to the original two systems. In developing business strategies, we never get far from the dollars.

Some of the equipment costs have been reduced over time with improved technology, but land, building and legal costs have increased as local communities fight proposals for new base stations. The result is that the cost of an average base station has remained at about \$1M over the more than two decades of cellular **voice** service. It has become more expensive in the last decade with the more complicated base stations of the smart phone era, but that's a story for a future lecture. The \$1M cost of a base station from 1985 to 2008 is confirmed by data published annually by the Cellular Telephone Industry Association (CTIA), as we see below. Why do you think it was increasing after 2008? Why might that matter?



The CTIA Data is extremely useful for our purposes in this course. It covers things like total investment, total revenues, number of cells, number of users, and number of “minutes” of talking that we will find very useful. We have gathered the CTIA data over the 35-year history of cellular onto an Excel spreadsheet and used the original data to calculate and plot certain other important information such as investment/base station (which we show above), users/cell, minutes/ user, revenues/ user and charges/minute. The visual display of the cost of a base station over time shows the value of such displays in understanding the information. Students should become familiar with the CTIA data, which has been added to the module **“CTIA Data and Other Related Files”** on Canvas. It will be used in the homework for this lecture and will continue to be used throughout the course.


Having had a quick glimpse at the future, we will now return to that period in the 1970s when Bell Labs was designing the system. The million-dollar cell would create economic and logistic problems involving how a cellular system would need to grow.

The impact of the million-dollar cell on system growth

Given that cells cost a million dollars, it becomes clear that we can't build a new system out of very small cells. A 3000-square mile system of 1-mile cells would have tremendous capacity but a startup cost of a billion dollars. It wouldn't achieve that cost of \$1000 per user until it had a million users. To reduce the initial cost, we would want a start-up system to have the largest possible cells. Based on radio propagation studies by Bell Labs in the early 1970s and the limited power of the cellphone transmitter, the 900 MHz frequencies being proposed for cellular allowed a maximum cell radius of about 4-5 miles. Based on 5-mile cells (75 square miles of coverage) a start-up system covering 1500-3000 square miles would only have 20-40 cells and cost \$20-40 million (far more attractive than \$1billion). Sadly, the inexpensive startup system doesn't have much capacity (not much channel-reuse). Even with an allocation of hundreds of channels our system with 5-mile cells will run out of capacity at (say) 20,000-40,000 users. We need an efficient way to grow gradually from a large-cell system serving 20,000-40,000 users to a small-cell system with millions of users, and we need to preserve all the older base stations that cost \$1M to create (slide 5, below).

Growing the system: We need to start with large cells to limit the investment per user

How do we get from to




While preserving all the old base stations that we paid a million dollars for

This leads to the process we call “cell splitting,” in which we put new base stations [the red dots in slide 6, below] midway between the old base stations [the black dots]. This results in a new cell grid with cells of half the radius ($1/4$ the area), and all the old base stations are still used in the new grid.

We will use the that same 7-channel reuse pattern with these smaller cells that we did with the larger cells, so the smaller cells will have the same number of users that the larger cells had, but in a quarter of the area. The system would grow through a succession of decreasing cell sizes, each with $1/4$ the area if the previous cells. Each cell-splitting would create a 4x increase in system capacity and a 4x increase in cost, so the investment cost per user would remain at about \$1000 (once it reached that early capacity of 20,000-40,000 users).

“Cell splitting” when more capacity is needed

- Preserving the old base stations ●
- New ones in between ●
- Each new cell will eventually have:
 - The same number of channels
 - In $1/4$ th the area
- So each “cell split” eventually yields a 4x increase in system capacity



Copyright 2009: Richard H. Frenkiel

6


A significant problem with cell-splitting

While this seems to imply a more gradual increase in the number of cells, cellular system planners realized that there would still be sudden and significant increases in the number of cells when a new cell size appeared. This would cause large and sudden spikes in system cost and the rather dramatic logistic problem of bringing many new base stations on-line simultaneously, with significant modifications to the old base stations (a different number of radios on different channels).

This problem is illustrated in the slide 7, below. On the left we see an initial cell grid that has reached “saturation” (it has used all its channels in the central 7 cells, and now must begin cell-splitting. On the right we see the result of conventional cell splitting. 30 new base stations are required (simultaneously!) to create a new smaller-cell grid that replaces the larger cells in the center of the system. Finding the land and getting local approvals is an unpredictable and time-consuming process. If some of the new base stations aren’t ready on the planned cutover date, there will be holes in coverage. Even if all the new base stations are ready on time, the number of base stations in the system will have more than doubled, at

a cost of \$30 million, while the number of users will increase much more slowly. As we will see, this huge sudden expense means that profit will be reduced, perhaps for years. Creating and equipping all those cells at the same time will tax the company's financial and logistic resources. Moreover, the new smaller cells will need different channels from those in the surrounding larger cells (the reuse distances will be different). The new smaller cells will each need about 16 channels to serve a small increase over the capacity that was served by the 60 channels in the larger cells they are replacing. The large surrounding cells will each need about 41 channels to allow a similar small increase. Thus, the configuration on the right is using almost all the channels ($7 \times 16 + 7 \times 41 = 399$). Another small increase in capacity will require the remaining large cells to split, requiring yet another large number of new base stations. During system development in 1976, this was seen as a hugely expensive and logistically daunting process with no obvious solution, even though the system development was basically complete (a service trial with 10 cells and 1000 paying users was planned for 1978).


Replacing the larger cells with smaller cells would be expensive and logistically challenging



Many new base stations would have to be ready at the same moment.
They would use at least $\frac{1}{4}$ of the channels, thus causing more splitting.
(surrounding cells would have to split very soon)

A solution to the cell-splitting problem


The solution, which turned out to be rather simple, is illustrated in slide 8, below. Instead of seeing the coverage in the central area as being transferred from the old large-cell grid to the new small-cell grid, it could be seen as coverage being shared by an underlying grid of large cells and a gradually emerging grid of small cells. In the beginning of the cell-split, almost all calls would still be served by the original large cells, with a very small amount of new capacity provided by a few isolated small cells as shown in red, each with only a single channel. In those red areas, users can be served by either the overlaid small cell or the underlaid large cell. Those three new channels are still used in the underlaid large cells, so to prevent interference we require those old base stations to operate as both large cells with 59 channels and small cells with 1 channel. They still have 60 installed channels, but only 59 can be used anywhere in the large cell. One of those channels must be assigned to a user in the smaller blue area. As long as there are users in the blue and red areas that we can assign those lonely channels to, we have achieved a small increase in capacity with no change in the system hardware. We only need to adjust our locating software to recognize which users are in the blue and red areas.



Continuing the gradual growth process (slide 9, below) we might see three more new base stations, and the new base stations might now have 2 channels each. In 6 of the old base stations, those two channels are now treated as small-cell channels serving nearby users. This technique remains efficient as long as there are enough users in the blue areas to make use of the small cell channels. Eventually the new small cells will completely replace the old grid of large cells, but the cost is delayed and the logistic problems have been solved. Remarkably, the average cost of the system during growth has been reduced by about 50% (simply by delaying the huge investment in new cells). The physical equipment is unchanged. All that was required to implement this “invention” (which became AT&T’s most important patent for many years) was some modifications to the locating and handoff

software (assigning some channels according to different rules). It seems that in many cases, “invention” is just a matter of looking at something in a different way.

We gradually add red cells (and to use more channels in those cells) as the system grows. At some point the small-cell coverage is continuous in the central area, and we have some small cells in other areas




The addition of new base-stations is gradual and flexible 9

“Snapshots” of a growing system

Slides 10-13 complete this discussion of cellular growth with hypothetical but realistic “snapshots” of a growing cellular system.


A look at a 400-channel cellular system-
The largest cells just before they begin to split



15 cells
 \$15M cost
 15K users
 1000 users/cell
 \$1000/user


10

System expands; cell splitting begins




11

More expansion and more cell-splitting



12

Mature system- Expansion and smaller cells




13

Lifetime costs based on the number of channels and the reuse distance

We will complete this reading with a look at the cost per user for a system throughout its lifetime, and the effect of various changes. Slide 14, below, is for the “standard” system using 400 channels with a 7-cell repeat pattern. It is verified historically by the CTIA data in Slide 16. After the startup period, the investment per user is essentially constant, which is nice for the service provider, since it allows predictable charges and profits.

The investment per user for a 400-channel system with a 7-cell repeat pattern




Copyright 2009: Richard H. Frenkiel

14


The CTIA data (slide 15, below) confirms this steady cost of about \$1000 per user, as well as the higher costs in the early years before cell-splitting begins

Checking the CTIA Data Investment/user 1985-2008



The effect of having to use smaller cells at Startup (hilly terrain?) is shown in slide 17. It mostly just affects startup (before cell-splitting) but could have some longer term effects at the system boundaries.


The effect of starting with smaller cells



17

The effect of getting 800 channels, 400 channels, or 200 channels (which we have discussed) is shown in slide 18, below. It leads to investments per user of \$500, \$1000 and \$2000. AT&T also claimed that the monthly charges per user for those cases would be about \$30, \$50, and \$80, respectively, which we will calculate in lecture 6. We will see how companies use financial calculations to estimate profit (or loss) and determine how much users will have to pay for service. The simple accounting that we will learn is a powerful and necessary tool for strategists, and we will be using it frequently throughout the course.


The effect of getting 800, 400, or 200 channels



18

Finally, in slide 19 below, we see the effect of having 400 channels, but with different co-channel separations (a 4-cell reuse pattern, a 7-cell reuse pattern and a 12-cell reuse pattern). An increased reuse distance might be caused by more varied signal and interference (perhaps caused by difficult terrain) or by deficiencies in system design (recall the effect of not using directive antennas at the base station). Going from a 12-cell pattern to a 7-cell pattern is the same as getting 70% more channels, so it was an important improvement in system design. We can see that reducing reuse distance is like getting more channels, which we know is very important in terms of system cost.

The effect of having 400 channels with a 4-cell reuse pattern, a 7-cell reuse pattern, or a 12-cell reuse pattern



Copyright 2009: Richard H. Frenkiel

19

In Lecture 6, we will introduce “corporate accounting for dummies” (which turns out to be more useful for our purposes than “corporate accounting for accountants.” This kind of effect will occur anytime we are involved in strategy. It is more important to understand that having half the channels will double the cost than to have a computer tell us that the cost will go from \$4753/user to \$8377 per user.