

We added a few algorithmic flourishes such as dynamic programming to compute optimal node degrees to balance storage and time. In retrospect these embellishments were less useful in practice; but they probably were essential for publication at a serious conference.

Jon Turner, the thoughtful chair of my department, told us that he would pay \$10,000 dollars for a patent. In those days that was a big sum. But there was a codicil. I had go to companies to "sell" the patent.

That initially sounded daunting. Surely being a professor was sitting in my office, thinking deep thoughts. But I reluctantly agreed.

One irony of the patent submission process is that our SIGCOMM submission omitted several technical details. This justly provoked the referees' ire. Fortunately, the patent submission process finished quickly and we were able to rush the paper to SIGMETRICS.

By then Marcel Waldvogel had visited Washington University. With Jon Turner, we had a second patent on binary search on prefix lengths which was more intellectually satisfying.

We decided to bundle the patents together. I contacted industry to market the patents. Everyone seemed to be building routers in those days, and so many companies, small and large, came to us.

Andy Bechtolsheim, then at Cisco after founding Granite Systems, called after we sent the ideas to Cisco. He told me that the multibit trie idea had already been done at Cisco, but they would not interfere with our patents if we did not interfere with theirs. Binary search on prefix lengths was new, but less useful to Cisco. This was disconcerting but it was a good reminder that ideas are often invented concurrently.

We used a non-exclusive model for licensing which I think was a good idea. We had a very helpful licensing officer called Ed Fickensher. It took a little time to educate him that our patents were not exactly equivalent to a miracle drug, and had to be priced accordingly. Once that happened things moved quickly.

We learned an important lesson about technology transfer. One must convince both the highest decision makers as well as the engineers responsible for implementation. Without support from the former, there is no power to make decisions; without support from the latter, there is initial enthusiasm but it quickly subsides into silence.

For example, at one software company there was an issue implementing multicast. Fortunately, Cheenu worked hard and shipped some code modifications. This direct interaction was crucial and instructive.

3. NETSIFT AND CONGA

After research in worm detection with Sumeet Singh, Stefan Savage and Cristi Estan, we started a company called NetSift. The signature extraction patent was li-

censed by UCSD to the company. The university was less interested in licensing revenue than in a potential acquisition or an IPO. After NetSift was acquired by Cisco, the NetSift team worked diligently at Cisco to build a chip to embody the ideas. As far as we know, the chip was never used.

A classic physics problem asks for methods of estimating the height of a building given a stopwatch. After a number of clever and strenuous methods involving dropping objects and casting shadows, the most painless method is to go to the building supervisor and say, "I have a very nice stopwatch you might like. Could you tell me the height of your building?"

In the same vein, I recently worked on a problem with Tom Edsall at Cisco and my student Terry Lam. Later, Mohammed Alizadeh joined Cisco and made crucial breakthroughs. Tom used a variant of these ideas, reported in the Conga paper, in a Data Center router. If technology transfer is the goal, then perhaps joint research with industry is the simplest route.

4. CONCLUSION

The DRR experience suggests that "giving away" technology facilitates technology transfer while the IP Lookup experience suggests that obtaining patents and pricing them reasonably, makes companies value even small inventions. The apparent contradiction is resolved by realizing that there are many paths to technology transfer.

Starting a company allows an idea to be more fully developed in context, while working with standards bodies allows the community to gather around an idea.

Since no approach is guaranteed success, it seems best to enjoy each attempt as an adventure, with the benefit of at least learning about new, real problems on the frontiers of research.

I summarize potential lessons based on my limited experience:

1. *Being contrarian:* Research that goes against industry trends may win you an audience in industry.
2. *Simplicity:* Simplicity may not be valued in academia but is often essential in transferred technology.
3. *Licensing:* The IP lookup patent and licensing model may be dated and not easy to replicate today.
4. *Rewards:* Intellectual Property Rights and a price can be part of transfer, but greed is poisonous (Fred Baker).
4. *Influencers:* Do not disregard either engineers or managers as influencers.
5. *Companies:* Founding companies, despite acquisition, may *not* lead to a change in practice.
6. *Partnering:* A number of companies offer university research programs. If the IP issues can be navigated, these offer the most accessible path to technology transfer working directly with engineers designing products.