The Company

(More on Ken Kolence)

Margaret Greenberg

A little background information will help us to understand the environment that enabled Katch and Kolence to get their start. Oddly enough, that begins with the founding of Stanford University in Palo Alto. Leland Stanford grew up the son of an upstate New York farmer, went to law school and settled in Wisconsin. About the time of the California Gold Rush that drew A.A. Michelson's father to the gold fields, Mr. Stanford's beautiful home and extensive law library burned to the ground. Like many in that era, he headed west to California, but not as a miner! He, like Michelson's father, was a supplier to the prospectors and made a fortune as a rail magnate. Some refer to him as a robber baron. In time, he became governor and US Senator. He lost his only child, a sixteen year old boy, to disease while on a trip to Europe. As he and his wife became older, they decided to build a university and donated a huge amount of land. However, the will stipulated that none of the land should be sold. Stanford University was born in the early 1890s. And the first student enrolled was Herbert Hoover.

An early professor of engineering, Dr. Frederick Terman, was a graduate of an Eastern University, MIT, and returned to teach at Stanford. He became quite a booster for Palo Alto. In 1935, two of his favorite students graduated in engineering and, like many young people from the west coast, decided to go east to work and obtain further schooling at prestigious universities. In time, he coaxed them to come back to Palo Alto, where they started Hewlett Packard in the house built by the first mayor of the city. Their lab was in a 12 x 18 foot garage behind the house. Their first big paying gig was for audio oscilloscopes used by the sound engineers for Disney's *Fantasia.* As time marched on, they began other ventures. This was before 1940. Today CMGers know them mostly for IT devices.

Meanwhile, Dr. Terman kept his ear to the ground for more recruitment opportunities. He was always thinking of ways to keep graduates in California, but more often the lure of esteemed Eastern universities and big corporate headquarters drew them away and many never returned. Very annoying!

Then, while chipping in to solve the university's endowment fund shortfall, he figured it out. There was indeed a way to fix his dilemma. It happened that the Stanford University administration was collectively scratching their heads over how they could best address this serious problem. Stanford was land rich and cash poor in part due to Leland Stanford's stipulation that none of the land willed to the university could be sold. Engineers are used to workarounds and his was to use those 700 acres. Lease the land to corporations for their headquarters or additional facilities. That would keep the boys

at home by giving them a place to work and, those who stayed outbound, a reason to return home. The university would collect money for the leases and students had a reason to stick around for postgraduate work. There would be businesses to hire them locally. And so the **Stanford Industrial Park** (SIP) was initiated. The first company, Varian, signed up in 1951. To Varian management, this was possibly not such a good move; therefore, they hedged their bets by designing their building so that it could be used as a school. Varian is still there.

As soon as Dr. Terman heard that William Shockley[i] wanted to start a company, he reeled him into the Industrial Park. Dr. Shockley had an interesting[ii] history in the War Department in the 1940s. But his mission in SIP was to make semiconductors. He was said to have a difficult management style with the result that some of his engineers (the self-designated "traitorous eight") bolted to start Fairchild Semi-conductor, where Kolence and Katch were contracting when they started to solidify their corporate plans.

Over the previous ten years, there was plenty of accumulated investment money ready to be used by industrious lads wanting to break out of the corporate mold and form their own company.

Resuming the adventures of Kolence and Katch Once again Kolence was faced with the prospect of leaving his "home" in the Bay area and moving to the CDC home base in Minneapolis and its dreaded winter climate. He was fired up but definitely not ready to go. So he and Dave Katch began to initiate their long discussed plan to start up their business and become independent.

> "When we were together at North American Aviation, Dave [Katch] and I had agreed that we wanted someday to form a company and had agreed on the name Boole & Babbage. I had reserved that name at the state agency that gave out corporate names and had kept it in force over the years. [… in 1967 while still at CDC,] I consented to start up our own company with Dave if he could get a contract before we started. Dave had some contacts in Fairchild Semiconductors. They had wanted a software system to use with a chip tester. For exactly what purpose I no longer remember, but I think it had something to do with culling out bad chips. Dave told them we could do it and they gave us a contract. This happened sometime in June or early July 1967, so we found some inexpensive office space, quit CDC, and started to work. We had no money to incorporate so we settled on a partnership. We took the name K&K Associates[iii]. That way I could say the first K was for Kolence and Dave could say Katch was the first K.
>
> Dave had worked for me at CDC primarily as a person who represented the Analysis Department on Design Review Boards and helped to produce procedures and related material for the "Programmers Handbook". So he was interested in the design process methodology, but not the measurement products. I wanted to build sampling monitors for the IBM system 360 computers, but several things prevented me from doing that for a while. First of all, we didn't have any money except what we able to earn. So we really couldn't afford to hire people who could work on these

products. Secondly, IBM had promised to provide Operating Systems that could simultaneously manage several different programs running concurrently. Unfortunately these systems had not yet become reliable. Besides which, none of us knew the guts of these systems. All in all then, we really couldn't even consider starting to build the measurement software products. So we focused on consulting work, hoping to get together enough money eventually to start building and marketing the measurement products.

Well, once again Dave Katch came up with a source of more money: His insurance man was associated with several people who were interested in investing in start-up companies such as ours. In particular he introduced us to Franklin Pitcher (Pitch) Johnson, a fellow about my age or a little younger who had taken a FORTRAN course and done a little programming. He was also a Stanford MBA with a bit of venture capital to invest. We explained the notion and uses of a software product to measure software performance and identify where the greatest performance improvements could be made. He discussed the idea and potential with his group of investors, and they came up with an offer. Basically, they would give the company (not Dave and me) $50,000 for 90% of the stock and would also guarantee us a loan of $100,000 when we were ready to begin product development. We had to put in $5,000 apiece for our combined 10% of the stock. Since we were the first software product company in what was just beginning to be called Silicon Valley, we felt lucky to find any capital to fund development. Dave and I discussed the deal, and decided to accept the offer.

By today's recent software venture capital deals, one might think we got taken pretty badly. However, in those days money was worth a lot more. One could live pretty well and even buy a house on a salary of $1,000 a month. For example, my salary when I left CDC was about $1,250 per month. We could hire some very good people at that money and with a guaranteed loan that we were given. (I must admit though that I wish stock options had been available at the time.) In fact, we were able to complete the products with this amount of money, plus what David and one or two others of us were able to earn on some consulting jobs."

About a month before incorporation, they hired Gary Holtwick to help with some consulting. Dave Morley joined about the time of incorporation.

"We incorporated as Boole & Babbage, Inc. on October 1, 1967. If you recall, Dave built the CDC6600 version of the program sampling monitor. He also had extended its capability to monitor the overall utilization of various hardware components such as the CPU, disk drives and "channels', along with the equivalent for tape drives, etc. This type of information was not needed for most systems prior to the advent of IBM System 360 computers, but was very important for these new systems. I'll explain more on the importance this of this later. Sometime after the first of the year (1968), our contract with Fairchild was cancelled. It – the contract – had been sponsored by Bob Noyce, but when he left Fairchild his replacement decided to follow a different approach to resolving the problems we had been working on. This

was actually a good deal for us, because it freed me to put the product development effort into high gear."

Thus, Boole and Babbage became the first software product firm in Silicon Valley. Kolence mentioned that there may have been others who did contract programming, but not product development. The building where Noyce independently invented integrated circuits [while with Shockley Semi-Conductor] is within walking distance of their office. Therefore, Silicon Valley was definitely under development—just not only for computers specifically.

> "It wasn't until Larry Welke started up his ICP directories that I really even knew there were product-oriented software companies in most other parts of the country. I can't think of any other names at that point in time where they were building software products *per se* here in what became Silicon Valley. I guess I did know about ADR and a few de facto software product companies based in Southern California. But at Control Data, I would not have been interested in them. When the Boole & Babbage products came into being, Larry was already publishing his directories and was about ready to give out his "Million Dollar Sales Awards".

He discussed why he didn't use the development process that he had just written up for CDC:

> "First of all, we had experience in building these products, at least for the CDC product line. Also, the experiments on randomization of the sampling process didn't need to be repeated, so we knew the algorithms we could use there. Finally, the products were functionally divided into a sampling program and a report generation program. The sampling program was the same for both products except for where data was collected from in the OS code. The Report Generation program was likewise basically the same except for the data manipulation subroutine. All of these reasons were important because we really couldn't afford the time and money it would have taken to use the formal methods that were necessary for the large scale systems of CDC.
>
> I need to give the products names to discuss why we could develop them so quickly. The program measurement product was called the Problem Program Evaluator, or PPE. The hardware configuration utilization product was called the Configuration Utilization Evaluator, or CUE. For PPE we needed to know how to interrupt the computer using the interrupt timer, which was fortunately a standard feature on System 360 machines. Then we needed to know where the interrupt stored the memory location of the address of the instruction the machine would return to when the interrupt ended. And, finally, we needed to find a way to determine if the interrupt had actually been completed or not. With this information we could determine the location that would become randomly active, and be assured it would only be reported once per sampling action. The code modules that extracted this type of data were called the "data extractors" and the report generators were called the "data analyzers".

Essentially the same process was used for CUE, but now the locations being observed gave the busy or not busy status of the set of hardware of interest. From this information we could then process the data to show when component hardware sub-systems were busy or not and when subsystems overlapped one another in execution. Generally speaking this was sufficient to determine where to change code or hardware to affect the desired amount of performance increase.

So the architecture of the products was almost identical for PPE and CUE and was easily coded. The hard part was finding the required data within each version of each release of the IBM OS. That is to say, we were really only building fairly simple utility products, not full-scale applications. We had built up a small development staff, mostly of ex-CDC people, and they worked together well. Dave Morley was the person responsible for the analyzer programs, and for specifying the data needed for the data analyzers. Gary Holtwick and Andy Chapman were responsible for the data extractors and the testing efforts. Andy was of particular value to us because he had been involved with these OS versions since their original “beta” releases about a year earlier.

The earliest users were three test sites in San Francisco. They provided us with the descriptions of the several different versions of each type of OS so we could then determine where within the OS we could find the proper sample data. Once we had done that and coded in correct locations, we then needed to make a number of runs using a set of known, typical applications and sets of data. Then we would let the installation people try it on their own. Assuming everything seemed stable, we would keep that version. However, there were so many different OS versions of each type being released in that time period that we were hard pressed to keep up. There were three different 360 OS systems promised for delivery, two of which were being tested and released at that time:

> **PCP** (Primary Control Partition) which could only accept one program at a time,
> **MFT** (Multiple Fixed Tasks) which could only operate over a given set of programs (I think).
> The last OS promised, but not yet in test at that time was called **MVT** (Multiple Variable Tasks).

As I recall, we had three test sites that we got to use in return for the software and some reasonable help in using them. Chevron was the main one, the City of San Francisco was a second site, and I can’t remember the last one. But the first commercial customer, one that we got money from, was the Guelph University Data Processing center in Canada. I can’t recall the exact order of the other early sales, but two others in San Francisco were the Bank of America and Del Monte. Two others I recall were AT&T in New Jersey and John Deere in Moline, Illinois.

We made the Guelph sale in January of 1968, and increased sales every month until we were selling about eight per month by about September of 1968. Further, our sales were all from word of mouth references and some SHARE and Guide presentations later in the year.

There were **no guidelines for pricing products**, and certainly none for the type of products we were offering. Plus, there were many people who honestly felt that they shouldn't have to pay for software. And, worst of all, there were **no budgets for paying for non-IBM software**. What we really wanted to do was sell PPE and CUE as a package, so that the systems people would use CUE and the Applications people would use PPE. Now in those days you could buy a disk – a 1311 model as I recall – for a little over $12,000. So we priced the two as a package at $12,500 so they could trade the purchase of a small disk for both PPE and CUE. Since we had to keep upgrading them both to work on the latest releases of each IBM OS, we also gave them a free year of "software maintenance" – that is PPE and CUE upgrades - and thereafter it was originally $500 per year. However, this fee was too low, and eventually we charged 10% of the purchase price. Back in those days the amount of money we are talking about seems ridiculously small now, but as I have mentioned that was before all kinds of inflation that hit during the late '60s and '70s.

We did hire a salesman, but selling software was such a new thing that he wasn't really able to bring in much business. As I said earlier, most of our business came about by word of mouth and also articles in Computer World describing the results of using the products. I went to SHARE and made presentations. Since I still had a number of good friends there I was able to set up sales presentations with them. But the best early leads came from a booth we had at the Fall Joint Computer Conference held in San Francisco in early November 1968.

Many, many people stopped by our booth, including IBM salesmen. Our products provided a completely new kind of capability to improve systems performance, both software and hardware. Since people had never seen what could be done using our techniques and products, they were naturally somewhat skeptical that their installations could be improved so easily. So, when people did get interested we made them an offer that few refused. It was a great closer. We would say, "Give us a signed contract for both products and we will come out on our [own] money. Give us three programs to analyze and give us 24 hours to fix them with help from your people who know what the code has to do. Then, we'll re-run them and if we don't save you thirty percent run time, we'll forget the sale." And at the same time we would run CUE and show them where their bottlenecks were, so they could also figure out where their system needed tuning. I don't recall ever having lost a sale doing that."

Most often they beat the 30% expectation. The few occasions when it wasn't met were so close the customer was satisfied that Boole and Babbage could provide a good return on their investment. The largest savings that Kolence recalled was

> "at Lockheed Sunnyvale, just down the road from here. They had a data analysis program that they used on telemetry data and it was running much slower than they had expected. They had already looked at the code and hadn't found any obvious reason for it. So, we ran PPE against it and immediately found the problem: there was a data correction routine that being entered for each piece of telemetry data, regardless if there was a problem in the data or not. As I recall, correcting that and a few other pieces of the code resulted in an **over eighty percent improvement in run time.**
>
> As it happened, I found out then that my professor from the University of Illinois, Dr Jack Nash, was head of Engineering Computing for Lockheed. I didn't know he was in the Bay Area, and was delighted to see him. After speaking with Pitch Johnson, who was Chairman of our Board, I asked Dr Nash to become a Director of Boole & Babbage. He was most pleased to be on our Board, and was very helpful in many ways."

Kolence was concerned about IBM reacting to the Boole and Babbage programs being run on their computers. He hoped to eliminate the possibility of their concerns by excluding IBM's system software, compilers and sorters from evaluation. If there were issues, IBM would be under intense pressure from their customers. Another thing to remember was the IBM sales staff pushed product to solve performance issues. And the salesman often was very convincing. In other words, the Boole and Babbage software was messing with the livelihoods of the salesmen. From his perspective:

> "It would have been very difficult for IBM to provide software that was efficient over all the workloads various customers might have. Remember the problem that CDC had just trying to get the same kind of efficiencies for just a small number of customer print workloads. Actually, PPE could even have been built to provide histograms of even the OS in execution. I figured that the last thing I wanted was to give IBM a reason for blocking PPE from working. So we put code into the PPE data extractor to prevent it from collecting histogram data from IBM products. Well, I was probably smart to have done that with PPE, but as it turned out the real threat to IBM was from CUE and I didn't realize that would be the case.
>
> Sometime around 1970, the Supreme Court ruled that the Patent Office had to give out patents on software products, so we prepared a file application and submitted it. The two inventors given in the patent were Gary Holtwick and myself. I felt Gary deserved to be my co-inventor since he had done the programming on the extractor. (Thinking back, I should have also listed Don Lytle and Dave Morley for the work they did at CDC.) The US patent Number

3,644,936 was awarded on February 22. 1972. We were unable to enforce it however for several reasons, one of them being that the US government claimed that their contracts with us gave them rights to use the methods in our patent. It was a big mistake to accept that argument, but our lawyer talked us into it.

While I was still at the Boole, there was only one new type of product that we developed, because of how, in my opinion, IBM tried to put us out of business. This was a product to try to reduce the amount of seek time on disks, and thus I/O wait time, by reorganizing the placement of these data sets. This product, whose name was something like Disk Data Set Optimizer, had a data extractor that collected information which gave us the ability to calculate how much I/O search time could be reduced by moving workloads to different locations on a given disk or to other disks. The original thought was to use Linear Programming techniques to determine an optimal data set placement from the data collected. However, this proved to be extremely difficult to do. Instead we used a heuristic technique that consistently produced very good results. **Most, if not all, of the disk defragmenter software products today ultimately descend from this product.**

It was an excellent seller and helped us to weather the difficult times that were coming up for the Boole. One of its most important results was that it got rid of the typical Data Center practice of assigning all data sets used by a department to its own 2311 disk pack[iv]. That step alone vastly improved performance in shops that had used this technique.

First of all, I have no proof that IBM as a company did what I think they did or why they did it. So this is my personal view of both things. Regardless, the events I will describe did happen. Let's start with the reason I believe led to everything. Back in the '60s and early 70's, companies normally asked their IBM reps to help them with the equipment plans for the companies. **I understand that IBM established levels of what was called "account control", the highest level being when the IBM rep was the de facto person who decided what equipment the company would obtain. Once the time sharing System 360 came into operation with either the MFT or MVT operating systems, the problem of determining what was the right equipment to get, became impossible to solve by the previous methods.** This was because when you are executing a mix of software on a time sharing basis, the bottlenecks that occur are a characteristic of that particular workload mix. A different mix of programs and you may well get different bottlenecks. The problem is that **without some measurement tool by which to see what is happening with different mixes, one has no idea of how to reduce bottlenecks and thus improve throughput times.**

Apparently, the IBM reps to most companies suddenly were overwhelmed by complaints by their customers because the equipment they had purchased was not giving them the throughput they had thought they would get. This certainly

happened at most of the Northern California and San Francisco area sites. When the CUE product was brought in and used consistently, the company technical people found they could figure out themselves how to resolve their throughput problems. The companies began to establish their own equipment planning groups, since they could get the data whenever they needed it. In turn, this meant that the IBM account rep's superiors were not happy. Eventually, this was happening in a large enough section of their customer base that they decided to do something.

That something, according to the sworn testimony of a person (whom I shall refer to as "Mr. J.") directly involved in the effort, was to build a copy of the two Boole & Babbage products and give them away free to their customers. I believe they expected this would cause the Boole to go under fairly rapidly. They were almost right, but they forgot how much maintenance effort it took to keep those products reliably working. The account that was selected to produce these copies was that of the Stanford Linear Accelerator or SLAC in the hills above Palo Alto. SLAC was (and still is) a world class physics center concerned with experimentally investigating and verifying various predictions of particle physics theories. I believe that SLAC administrators were unaware of what was happening, but that this was not true of several people in the computing facility. To begin with, we were invited by them to make a presentation on PPE and CUE, how they worked and how one could use them to solve various problems. This we did, expecting them to let us know if we could give them our usual offer to go there and run our products, collect data, and show them that our claims were true. We did get a few calls, and when we called to see if they were ready to buy we got vague promises but no action. So we finally gave up on trying to make a sale. Somewhat later we received, from some unknown but honest person, a copy of an internal SLAC computing center memo that outlined the progress two IBM sales support technical people had made toward building copies of PPE and CUE. In it, they said that it could be released and distributed to COSMIC[v] within a fairly short time. COSMIC was and may still be a NASA sponsored system for the distribution of free software to government agencies and whomever else might be interested.

To say the least, I was shocked, angry, and most of all scared at what might happen to our sales, which were doing quite well. My first inclination was to talk directly with a person who I assumed would be able and willing to prevent their copy from being distributed free. After discussing this with Pitch and others, it was agreed that was the best thing to do. The person was at that time a highly respected officer of the National ACM organization, or had recently been, so I naturally assumed he was a reasonable and ethical person. I called this person up, told him I understood SLAC was getting ready to distribute their copies of our products, and wanted to speak personally with him. He agreed to meet with me the next day. I went up and he said, 'Let's take a walk.' We went outside and walked to a place where no one was around. He admitted that they had built the copies of our products and that he was going to release them to COSMIC. In

other words, he was the fellow in SLAC that had responsibility for the project. I remember saying to him: 'Please don't release this to COSMIC. If you want to make it for yourself that's your business, but you're trying to kill off this whole thing. We are the ones that invented it. We are the ones that invested our money and our effort to make sure it would work and to provide it to other people and to keep it maintained'. And he looked at me and he said, 'You know, what you should have done is made it in hardware, not software'. That made me very angry, because he had to know that there was no way either product could be easily put on a board. I must have shown my anger, because he looked right at me and said. 'I am going to distribute it' and walked away.

I have never been so angry with anyone before or since. Here is this guy, a big professional, you know, and he's doing his best to help IBM keep in control of all the installations, configurations, sales stuff. From that point on I had a very negative attitude toward an awful lot of people in the ACM. That attitude is gone fortunately. But his decision certainly changed my career.

Well, anyway, SLAC went ahead and released their software to COSMIC. But it was a strange thing that the day that it became available from COSMIC, in every account that we were in sales negotiations with the IBM rep came up and said something to the boss like this: 'Look we've got this software in COSMIC now, it's free, it does the same thing as the Boole & Babbage software, and you don't have to buy it.' All of our sales activities stopped. We were dead in our tracks. The only thing we were selling was our new disk data set organizer. Fortunately, we had our consulting work too. This went on for several months. But we were barely making it and I, in particular, was going around just working my tail off.

To top things off, a little while later I went to a GUIDE conference and tried to drum up some sales. And here's this guy, Mr. J., talking about the measurement products they had built at SLAC. He never once mentioned Boole & Babbage but he really strongly intimated that they had invented it first, and that 'someone else' had stolen the ideas and commercialized them. Apparently he didn't know who I was, because although I knew his name, I had never met him before and here he was busy trying to take credit for inventing all these products. I lost it. I lost it completely. I really lost it. And I got up and I told him what I thought he was, what kind of a person he was, et cetera. And I did this in front of a crowd of people. Well, I was not too popular with the GUIDE people after that, because they believed the SLAC/BM's story. So that outburst got back to Boole & Babbage, to Pitch certainly. And I lost my temper on some other occasions. I had been working so hard to generate sales that I couldn't remember if I had said something, read it, dreamt it, or what. I began acting irrationally. I was very close to a nervous breakdown. I blew my stack at Gary Holtwick, the guy who wrote the original extractor code. Gary went to Pitch Johnson and complained about me - I don't blame him – I had lost my cool completely. There was no question about it. So I guess Pitch decided he had to get rid of me, but first he had to find someone to take my place and get up to speed on what we were doing.

I could tell that I was being set up to be fired, since I was being forced to fire people like Dave, Jen Brian (head of marketing for the Boole) and other top people I had brought in. There were twenty-five or thirty employees when the COSMIC release occurred, split about evenly between consulting and support teams for the products. There were considerably fewer afterwards. David Katch, Jen Brian, and myself were really all from North American Aviation through Control Data. But Gary Holtwick[vi], Dave Morley, and several others were from CDC.

I turned in my resignation every month for about 6 months and Pitch would refuse to accept it. Finally, after he had found his new president and had him up to speed, Pitch fired me. I don't recall exactly when this happened, but I'm pretty sure it was in the summer of 1972.

Quite frankly, I was relieved the entire thing was over for me. I went home and slept for a month or so. But I never would have fallen apart the way I did had SLAC and IBM not done their dirty work. Looking back at things now, I also think the field of computer measurement would be far advanced from where it is now."

A rather sad series of events. Two guys had a dream to start their own company, build a great set of products and ended up losing the company by being fired from it. Some other strong development people also had to be let go for "guilt by association". Nevertheless, the company lived on until 1998 when it was bought by BMC[vii]. The products survive in **Mainview[viii].**

Next time, the CMG connection.

---

[i] One of three recipients of the Nobel Prize for invention of the transistor.

[ii] He was requested to prepare an estimate of the casualties for the impending US invasion of Japan. It is this estimate that convinced President Truman to use bombs rather than risk soldiers lives. Read about at http://www.endusmilitarism.org/casualty_estimates-giangreco_rebuttal.html

[iii] One source noted that Pitch Johnson objected to the name of K & K Associates because of a possible confusion with the KKK.

[iv] This practice reappeared in distributed environments where each department or group insisted on having its own server (s) that no one else could use.

[v] http://www.openchannelfoundation.org/cosmic/

[vi] Remember that one of the reasons Pitch Johnson had for firing Kolence was Hardwick's complaint.

[vii] This boosted BMC, which had just been downgraded and resulted in a 14% drop: Boole and Babbage stock was earning $.34 per share.

[viii] BMC product