

# Recollections of Tymshare

Mark Kahrs

## In the beginning...

In 9th grade (1966), my Palo Alto high school (Cubberley) announced a new option called the “Research Observers Program” (ROP). The idea was to place students in research labs so that they might experience the joys and triumphs of research (some boredom too I suspect).

At that time I was entranced with the HP 8410 Network Analyzer and so the school managed to match me up with the passive components group in the Microwave Lab at HP on Page Mill Road. I would ride my bicycle up to HP and hang out — my education being insufficient to do much of anything. At some point, it was mentioned to me that I might want to learn how to program computers. They had a room with ASR-33 teletypes connected to commercial timesharing firms. I began by getting a book by McCracken on BASIC from the HP Library and started using the GE Timesharing BASIC. After I had that figured out, I moved onto the Tymshare account and started learning assembler (I had heard that was the real stuff - closer to the machine). My nerd buddies at Cubberley (Hal Sampson, Bill Parrish) had meanwhile used the ROP program to go to Tymshare. Bill Parrish tells me that it was Arden Scott who had the idea of employing high school students. The “Tech Division” headquarters was on East Meadow Drive and across the street was the machine room and a room full of teletypes. This had an additional advantage it was 10 minutes by bicycle from my house. So, on the weekends, I would head on over to the common teletype room and burn up HP money (I had no idea how expensive it was — in retrospect, I probably burned through thousands of dollars, if not tens of thousands of dollars). I had figured out the HP account scheme, so I was spreading the joy around to different accounts as I worked on my assembly program. My assembly program was just a vehicle to explore different system calls and I still remember using the relabeling (remapping) system call — which promptly crashed my program. Eventually, HP did find out it was me since I had put my name in a comment and I was gently but firmly yelled at...

## East Meadow Drive

But it was in that terminal room that I first met Norm Hardy. I also met Ann Hardy there and others, esp. the machine operators (Syd Reid comes to mind). I managed to crash the machine a few times and that made me a bit notorious with the operators (imagine trying to use the tape system call with inadequate protection).

Later in the year, my mother insisted that I ask Norm about a summer job. Being all of 16, I was terrified of Norm, but eventually picked up enough courage to ask him. And he said he’d look into it. What he found was the same job that Bill Parrish had, which was writing programs in SUPER-BASIC (Tymshare’s enhanced BASIC) to ensure that the programmers responsible for enhancing SUPER-BASIC were not introducing bugs while fixing others. I don’t recall ever finding a bug with my programs. I was paid the glorious sum of \$1.50 per hour, which was well below minimum wage (I was a consultant of course, which meant that

Tymshare didn't have to pay benefits). I deposited all the money in the bank for my future college expenses.

The "Tech Division" office was across East Meadow drive (from the machine room) in a small non-descript building. As I recall, it had two hallways connected by a common reception area. Nearby was the first Xerox copying machine, the model 914. I remember how amazing it was to get a copy of a piece of paper without photography. Being in the "Tech Division" was like a large family; it was an exciting time. I met the other systems types: Verne Van Vleer maintained the Executive (what Unix would call a shell). Verne had a private plane and once gave me a ride on a weekend. Guy Blood was responsible for security and accounting and frankly he scared me. He had a volcanic temper and was a chain smoker. Eventually he got used to me and he wasn't so scary. Walt Main and Steve Saperstein were on the computer languages team. And Norm and Ann Hardy and Dave (surfer) Gardner were working on the Operating System.

In my junior year (1968), I continued to work for Tymshare. I was now paid \$2.10 an hour (a 40% increase). By this time I was driving my parents a little crazy by going directly to Tymshare after school and missing dinner. I would also bring home assembly listings and spend the evening trying to figure out how the monitor (operating system) worked (and not driving fast cars and chasing girls). If I had any questions, I would ask Norm, who would calmly explain what was going on. I didn't realize it, but the entire monitor was being rewritten. I still recall Norm explaining to me why the longer 6 line page clearing code segment was faster than the 3 line segment. The biggest change was the introduction of the "Random File System". This permitted character level access to the file system. Today of course it's a no brainer — in Unix it's just lseek. But during a time when block access was all you had, this was a great innovation. The changes included adding new instructions ("pseudo-ops") and system calls.

One of my (now) funniest memories was attending a seminar where Feldman's semantic language paper in the CACM was discussed. The first question was "Does anyone here not know BNF?". I raised my hand... and no one else did. Well, I was only 16 at the time... Tymshare was quite interested in programming languages since that provided the "value" to the system. In fact, the original Berkeley software was deprecated — while CAL (Lampson's version of JOSS) was available, nobody used it. The money was in SuperBasic.

My entire senior year (1969-1970) was spent in the following manner: I would do school homework during lunch and work at Tymshare immediately after school and at night.

Occasionally I would bring home a ASR-33 and a modem (110 baud only, these were acoustical couplers). It was a novelty but it tied up the phone line thereby annoying my parents again — they had patience. I was quite interested in the Berkeley Project Genie subsystems and I wanted to bring over various languages. I brought the new Berkeley one-pass assembler (NARP) over and although slightly incompatible with the existing assembler ARPAS it was a lot faster. Considering the speed of the machines, this was an important improvement. I was interested in everything Berkeley had to offer, including LISP, but the different disk system made it a difficult job to do the port to the Tymshare system. I also kibbitzed on how Tymshare should interface to the Comshare customers we acquired. I worked with Dave Gardner on trying to add "subsystems" to the existing list (subsystems were the languages and processors of the system) but when we tried, it broke the accounting

system. That was definitely a no-no.

Because I was so deep into the system, I was permitted to have system privileges. These were assembled into the executive and I recall the excitement I had when I first used my new exalted status. I still use that password. I added a feature to DDT and also wrote various system utilities. My best system program was a version of systat for the 940. I never did tell anyone about it (again, more's the pity) but it used all of my knowledge about the internal system structures and system calls.

I managed to crash the machine once because there was a new system call that permitted disk reads without waits. But if you got in a loop, you could fill the monitor's command buffer. Dave Gardner fixed that one telling me that even with my exalted status, I still shouldn't be able to crash the system.

With my studying of the monitor listings I soon became a walking system call reference - I was often stopped in the hallway to ask "What do you put in the registers for BRS ..." (BRS is "BRanch to System" - the system call instruction). Ann Hardy then assigned me the task of maintaining the system call manual. I was able to use my knowledge of the monitor to rewrite the various pages of the manual. I would do this for the next 3 years. This included rewriting the entries into lower case (hard to imagine a life without lower case, isn't it?).

But soon enough my idyll would be over: Tymshare would move to Cupertino circa 1969.

## Bubb Road

The new building on Bubb road was frankly one of the ugliest buildings I have ever seen. It has plenty of room for machines but precious few windows. And rocks on the outside walls. It was hideous. And, I had to borrow the car to drive there.

The first summer after they moved, I commuted with a secretary (Jeri Crockett). I would ride my bicycle to her house and then drive in with her (I agreed to split the cost to save her money). In later years, I would ride in with Steve Saperstein and Walt Main. More than once they would show up and I was still asleep, having forgotten to set the alarm.

The new building had two partitions: one for people and one for machines. A single door connected the two. When the lights blinked, there was a mad rush for the machine room because the machines had to be rebooted (manually) as fast as possible. Every second down was lost revenue. Clearing memory on the 940 was designed like a nuclear weapon: you had to push two buttons at once. When the systems were being rebooted if you raised one of the switches, the memory ("core") would be written to a spot on the disk. You can then read that spot when the system came up and copy it to a file, where it could be inspected with DDT. I loved to do this and try and find the bugs - I'm not sure I ever did. I recall one difficult monitor bug when the disk would "lose" and interrupt and eventually the system would crash. More than once I helped reboot machines over the weekend.

It was at Cupertino that Tymshare designed the channel interface to the 940 so that the new IBM 2314 disk drive (with new state of the art disk packs!) could be connected. Also, it was at Cupertino that Tymnet would be constructed. My office was in the very center and I shared it with two other youngsters: Mike Geary, a recent dropout from CalTech (who

arrived with the Sigma 7) and Gary Gere, who found his way there somehow someway. Our office was in the very center of the building. Later Gary Morgenthaler would join us.

Before the IBM disks there were Data Products disks. These were at least 3 feet in diameter. I remember watching one repair job. They placed a cloth on the floor. Then, using calipers, they removed the disks one by one off the spindle and placed them on the cloth in a stack. The procedure was reversed to reassemble. The final step was to vacuum the surfaces through a fine cloth. Given today's disks, it's hard to imagine. Some machines had fast disks and others had drums (still!) for swapping.

At the end of the summer after I graduated from high school, it was time for me to go to college. I was given a grand send off at a fancy restaurant in Los Gatos (complete with fashion show) and a copy of "Fundamentals of Algorithms" by Knuth with my name written in the front cover by Ann. I still have it and use it.

In the following two summers, I returned and continued my renewal of the BRS manual. Meanwhile, XDS had demanded that Tymshare experience a Sigma 7. I remember my disappointment when we tested the system; clearly XDS had not learned anything from the Berkeley modifications to the 930 and the underlying software concepts. Eventually, it was shipped back to XDS. As is well known, the Sigma series was a failure and XDS became subsumed into Xerox. They did however build the Xerox Alto, another favorite machine of mine, but that's another story.

In place of the Sigma 7, Tymshare ordered DEC PDP-10s. These proved to be popular with everyone and soon there was a whole host of experienced programmers (Bill Weiher, RPG - Richard P. Gruen). They set about connecting it to Tymnet and interfacing it with the accounting. Another priority was to create a command interface that was more like the 940 (TYMCOM-X) so customers could migrate to the PDP-10 with minimal disruption. I wrote a 940 to PDP-10 assembly language translator but never seriously tried it out. A pity. I also assisted in the connection of the line printer for the first PDP-10 but crawling underneath the raised floor with an adjustable wrench to undo a nut on the bottom of the pallet. There was no way Richard Gruen was going to fit.

Walt Main and Steve Saperstein wrote a compiler-compiler called SIMPL that could be used for language implementation. Mike Geary and I started to implement a version of APL using SIMPLE but quickly ran out of gas.

My final summer at Tymshare, I worked on a parser for a database language. I was still besotted by Floyd-Evans parsers and so implemented a version of a report from the University of Illinois. I had to explain this to Carolyn Diehl (who changed her last name to Rose) before I left - I doubt it was ever used. I would later implement a much better Floyd-Evans parser for my M.S. thesis at Berkeley (on the Alto).

## Afterward...

After I graduated from college, I obtained a grunt programming job at Xerox PARC. It wasn't much fun but Xerox PARC itself was an amazing place. But that ended my time at Tymshare. My job at Tymshare was formative in so many ways:

- I found that being nerdish wasn't unusual in this environment, in fact, it was almost expected and valued. So maybe it was OK after all.
- I was encouraged to learn the 940 system and the assembler. This background has been beneficial whenever I look at any program or system to this day.
- Tymshare literally paid for my college education: my summer money would pay for the following year. When the last quarter was over, my bank account would be empty. (I admit, the UC system was very cheap).

I remain indebted to Tymshare and the staff who put up with me during a very important time of my life. Without Tymshare — particularly Ann and Norm Hardy, I would not have been encouraged to pursue computing.