



**A computer timesharing application at Fairchild  
Semiconductor, Inc.**

by:  
Mike Humphries

CHM Reference number: X9406.2021

© 2020 Computer History Museum

“The Cloud” has slightly different meanings for different people and in different contexts. But in most cases it is the term for the rich and powerful environment providing access to applications, content and data from remotely located resources. From the user perspective, it’s just there and available, accessible from almost any device you can name from desktop, portable, or mobile. It is powerful, convenient, cost effective and releases users from physical constraints in performing their work.

Some have noted how today’s “Cloud” can be viewed as a more evolved and ubiquitous form of another technology-enabled environment that existed in the mid-1970s -- more than fifty years ago. It was called *Computer Timesharing*. This is the story of how and why it was used through a personal application example. Timesharing changed many aspects of computing from expanding the use of digital networks to remote, online customer support -- pioneering things we take for granted today.

This account may give you an idea of the kind of real industry requirements in the 1970s that made using a computer timesharing service an alternative to in-house corporate computing services the typical large company offered its own internal organizations. Let me use an example of the business requirements that drove customers to use computer timesharing by recalling my own involvement in this very specific real life 1970s computing application.

Here is an overview of the business problem that faced me:

1. Gather in-house sales data to generate an accurate report to an industry association by a monthly deadline.
2. The analysis in these reports had a significant payoff in making business decisions.
3. An interactive environment was needed for fast and regular completion of both tasks: gathering accurate data and doing timely analysis.
4. Reduce the turnaround time for sending and receiving reports, so that subsequent analysis was as timely as possible.

My own experience with computer timesharing started with a classic case of a departmental need to get some specific computing done on a recurring basis. In late 1971, Fairchild Semiconductor was a semiconductor technology, sales and marketing company that offered a huge number of product families and device variations. The industry was very competitive, so prices and market share were under constant scrutiny. Fairchild had been an early leader of the industry but was in danger of losing its competitiveness. Many critical sales and marketing strategies were affected by what these numbers indicated.

Fairchild’s Market Research and Planning group was located in Mountain View, California. They did research of their own to suggest trends and strategies but they were often asked to respond to questions from senior management as well as others in sales and marketing. Quick answers to some of those questions were very important, but not something easy to do in a lot of cases.

The biggest source of meaningful data was from a semiconductor industry association to which virtually all major semiconductor manufacturers belonged. On a monthly basis, each member company reported unit and dollar sales for each device type they produced. The device types were defined by the association so that reporting could be uniformly accomplished no matter what naming and numbering nomenclature was used by each member company.

The task each month for Fairchild was to first produce a report for the association of standard device types, units and dollars from an in-house report that covered each Fairchild proprietary device number. The first big time-consuming and error prone task was a lot of crunching of Fairchild device categories resulting in a paper report on a form supplied by the association.

These reports were sent to a representative of the association who held each companies' individual data in secrecy. The reports were consolidated into one big report that aggregated all the association members' sales by generic device type. There was no identification of any individual company's sales numbers -- just group totals. As a member, you received back a copy of this amalgamated report.

Analysis and in-house reporting could then begin. Since you knew what you had reported by generic device type and you now had last month's industry figures by generic device type, you could make all kinds of calculations that showed how Fairchild was doing relative to the competition. Standard monthly reports were produced around market share and price/unit trends but there were also a lot of reports requested on an *ad hoc* basis depending on the business climate at the time. Both of these processes were almost completely manual, using a group of clerks who used the in-house reports to construct the industry report, submit it, then after sitting on their hands for a few days, waiting for the return of the consolidated report to begin work to generate in-house market share and other reports as needed. Except for the monthly Fairchild sales report (stacks and stacks of fan-fold line printer output) produced by Fairchild's mainframes and Management Information Systems (MIS) department, the rest of the processes were manual, at best assisted by electric calculators. That is the way it was for virtually every device manufacturer.

My involvement began one day when I was at my very first job fresh out of Georgia Tech and working for Fairchild Semiconductor (and Mike Markkula.) The head of Market Research and Planning (MR&P) came into Digital Product Marketing where I worked in Medium Scale Digital Devices as a product marketer. He was looking for someone to help complete a computer automation task that was suspended in midstream. MR&P had already made a decision that the current manual system described above had to be changed drastically. They had found MIS could not do the interactive parts of the requirements at all, and the application would only be addressed sometime in the future when the backlog of other MIS projects allowed. In those times that was normal. New computer applications often took months before they could be assessed for requirements and design by MIS. And even longer for implementation. Plus ongoing support for operations and the inevitable changes and enhancements posed an even bigger problem.

To be fair, most companies' MIS groups were running as hard as they could to stay up with requirements for core company needs: manufacturing scheduling and reporting, sales reporting, payroll (although most came to rely on third party vendors for that) and other basic business needs. To get them to focus on a departmental need required some juice. And even then, the line of projects in front of you could be long. They had small staffs for operations and even smaller for evaluating and building new applications. Support for those applications was also sparse. The computers of the time were mainframes that required a special environment virtually all contained in a glass room with raised floors. Power and air conditioning had to be special. Access was as a matter of necessity limited. Remote access was essentially unheard of. Programs were scheduled like buses or airline flights and run as 'batch' jobs at mostly designated times. Applications of lesser importance were at the mercy of the more important applications deemed 'critical'. So you can see why there was an opportunity for a better answer and why people embraced them.

That was the reality of available choices for Fairchild's MR&P to decide how to meet company and department requirements.

Fairchild Market Research and Planning (MRP) had elected to write computer applications themselves for 1) association reporting commitments and 2) in-house data analysis and custom reports. There was an MRP department employee who chose a solution enabled by using a

computer timesharing service. He selected Tymshare as the service provider and Super Basic as the language in which it was to be written. The guy doing the project was only part way through implementation when he left the company. Nobody else knew much about what he was doing, and there were no design documents or other notes available. The MRP folks approached Product Marketing looking for someone to help them complete this project. I heard their requirements and was intrigued by the challenge. As a result, I signed on for what was to be a six-month commitment to get the project completed, resulting in it operating on an ongoing basis every month in the future.

At that point, my total experience with computing had been three courses I took at Georgia Tech as a part of obtaining my EE degree. In those courses I had programmed in ALGOL (a Fortran-like language.) I had no prior knowledge of Super Basic or any of the other resources Tymshare offered that were needed to get this done. This is where the simplicity of Super Basic, its self-documenting nature, and the excellent training and documentation from Tymshare, made up for my lack of knowledge.

I plunged in, took some classes from Tymshare in Super Basic and its OS command level Exec plus Editor, the interactive application for editing code. The Tymshare service model included teams of Tymshare field sales and support people that learned what each customer was doing and needed, and not only advised them on how to do it but often directly helped get customers through tough spots.

The first task was to complete the data entry and device type consolidation process so we could meet the reporting deadlines for the monthly association report, as well as reduce errors while dramatically cutting down the time to complete this task. We chose Datapoint CRT terminals with magnetic cassette tape for the data entry (our own MIS could not supply us with a computer readable version of the company raw sales report.) We took delivery of the very first ones delivered and immediately put them into use. The terminals and tapes allowed data entry, verification, and editing offline.

When the monthly data was all entered on tape, the Datapoint terminal was then connected to Tymshare's network and computer via modem, the data uploaded to our department Tymshare account, and then formed into sales data files. The first application I wrote processed the raw unit and dollars data given to us by MIS into product groups by generic device type. Next it formatted the device types and accompanying sales data to fit the report format we had to submit to the industry association. And finally, we produced a properly formatted report using a small printing terminal. Then the report was off to the association -- they were not able to accept an electronic computer file of the report.

When in two weeks or so we got a consolidated industry report back from the association, the next phase of the work began. The report received was entered into our Tymshare account disk files much like the company data had been: by typing it onto cassette tape and then uploading it. It then became a pretty simple task to use that consolidated industry member's data and our own data to produce a number of standard monthly reports and analytics that were a top priority for Fairchild management. Using matrix operations, we pretty much divided the Fairchild matrix (units and dollars by generic device type) by the consolidated industry matrix to produce market share reports by device type. By virtue of the previous month's data, which we stored in our account, we could also plot historic price and volume graphs for Fairchild and the industry.

We reduced this report/analyze/report cycle from almost exactly one month to about a week. Now the internal analysis reports were three weeks earlier than before, enabling timely analysis and decisions that could help drive the business. Three weeks sooner was critical. Tymshare and computer timesharing helped make us heroes!

The biggest resulting challenge was the ensuing demand by management for *ad hoc* reports and analysis. Now that they had data much sooner than before, they thought of a number of other cuts of the data they wanted. This was my introduction to what is the norm in computing: improve the delivery of data and the result is a request for more.

Depending on market conditions, success, failures, competitors and Fairchild's new product introductions and production problems, there were many new demands for reporting, which differed greatly each month. At that time (1971) Tymshare had no general report generation facility. Each report had to be written in Super Basic. Even though portions of previous report code could often be re-used along with new code, there was no easy or quick way to supply those newly requested reports. Even so it was much, much faster than the old manual method, so it was deemed acceptable.

Early versions of report generators would come along within two years but it still amazes me that neither I nor any of the folks at Tymshare I worked with saw that we could generalize some of the code we had already written to allow new reports to be produced faster each time they were needed, It was an obvious business opportunity we were blind to, but we were too happy with what we were already accomplishing for the task at hand.

Here's a summary of the critical pieces that Tymshare (and other timesharing vendor) products and services provided that allowed the initial goals in MRP to be accomplished:

1. The availability of an on-demand computing service to build, operate and deliver applications to end users.
2. Languages that were interactive so that building apps was quick, allowing testing and completion to occur without waits for periodic 'batch' runs.
3. A highly reliable computer network that allowed the app building, processing and reporting to be done from any location at work or home with a phone line – via a modem and a terminal. The incredible hours I put in to develop the programs and maintain them were sometimes only able to be done remotely from home or other locations.
4. Terminals to fit each need from off-line entry and editing, online uploading, verification, reports and analysis.
5. At our end-user site, help and advice from the team assigned by Tymshare to my account. These people did the same support for other customers and were often able use that experience to suggest better and more efficient ways to accomplish our computing needs.
6. A behind-the-scenes operations staff at Tymshare. They kept the computers and network running and available more or less 24/7. They did regular tape backup for all customer accounts so that system problems or human error would not have much effect. For massive data requirements they could also accept on-site data input in magnetic tape or punched card formats.
7. A recurring schedule of classes locally for each of the major languages, the Tymshare operating system, and subsidiary applications like Editor.
8. Superb easy to read documentation made for non-computing professionals. There were quick summary guides and reference manuals, all very readable and with clear examples. They differed night and day from what was being supplied by the mainframe vendors to their MIS customers.
9. Evolving technologies that added help to building and operating applications. The first of these was a built-in 'Sort' capability that allowed an application program to call it, point it to a data matrix for sorting and return a correctly sorted data set. Before that the sort function had to be written into each program! Availability of several other related vital functions followed rapidly-merge, select, delete and report. Program operation became faster and maintenance was easier.

The method of working described above -- in which tasks are accomplished remotely by accessing distantly located resources -- was re-christened 'the Cloud' in the mid-2010s. As shown in the example above, some of the Cloud's functionality and basic philosophy of computing time as a resource like electricity or water were prefigured decades earlier with the mainframe-based model of computer timesharing.