



The Origin Tymnet

by:
Mike Humphries

CHM Reference number: X9407.2021

© 2020 Computer History Museum

The Origins of Tymnet--Part I

Tymshare began providing timesharing service on the SDS 940 in September 1966. The machine came from Scientific Data Systems with "CTE equipment" that attached about 16 terminals or modems, each with an RS232 voltage mode plug. Until recently modems had the same shape plug. The only practical terminal then was the model 33 Teletype from Teletype Corporation, a subsidiary of AT&T. It ran at 110 baud yielding 10 characters per second. It was upper case only, heavy and noisy. It was also cheap and reliable.

During the first year or so, customers dealt with the phone company to reach the computer center in Palo Alto, California. Most were in toll free calling radius. In 1967 Tymshare established another computer center near Los Angeles and developed a largely disjoint set of customers there. It planned further expansion.

There were frequency division multiplexers on the market that would handle about 16 Teletype circuits over one voice grade line. Each user was assigned a frequency and full duplex leased voice lines would connect the multiplexers. The Company tried these but they were expensive and error prone. Minicomputers, costing several thousand dollars, were then becoming available and it seemed clear that they could do a variety of communication tricks which would soon pay their way in saving communications costs. I joined Tymshare in Q4 1967 from IBM and LaRoy Tymes joined in Q1 1968 from Lawrence Livermore labs where we had met in the early 1960s. Because of our experiences together there we were both eager to pursue this adventure. Howard Steadman, already at Tymshare, dealt in hardware and encouraged us to suggest innovative solutions to interfacing with modems. Dave Schmidt was vice president of Tymshare and understood the technical and economic advantages to Tymshare of these schemes. LaRoy wrote all of the early Tymnet code. He and I did the design.

At that time the phone company (AT&T) was considered a natural monopoly. They held that modems were the sole province of the Telco. Tymshare had already produced acoustic couplers that competed with the Telco's 103A data sets (data set = modem), but AT&T's lawyers were busy fighting what they considered to be more serious encroachments on their turf.

Initially we found a strange little computer called the SPC12. It had 4K of 12-bit words. Most significantly it had parallel input and output, 12 bits wide, which could be interfaced with modems with simple voltage converters. The software would prepare and time the signals and there was no hardware devoted to timing the bits for each communications line. Precursors to this plan had been done at Livermore and in military applications before that.

The first deployment used the SPC12 as a fixed time division multiplexor, providing full duplex service for 29 terminals, each running at 110 baud. Pairs of such multiplexers were connected with a voice grade line, leased from AT&T, and two AT&T 201 data sets at 2400 bits/sec. There was no error control. These were deployed for a short while. We found the SPC12 too small and had located another machine, a Varian Data Machines model that used 16 bit words, and which could be expanded beyond 4K words. This machine, and its direct descendants, served as the Tymnet workhorse for the next several years. Howard Steadman equipped the machine with parallel I/O, like the SPC12. We programmed this machine to do statistical multiplexing which took advantage of the fact that not all lines ran continuously. This allowed more users per leased line. We added backward error control on the multiplexed lines.

With this phase of Tymnet (not yet so named, however), pairs of nodes would be connected by a leased line. One node, the base, would sit next to the 940 computer, connected to it by many cables. This still required the CTE equipment. The other node would typically sit in some Tymshare sales office. That office could access only that one 940 computer.

Several developments then began. Ann Hardy began the modification to the 940 operating system to supplant most of the function of the local node, thus avoiding work both in the 940 host and in the 620 node. The 940 and 620 exchanged teletype traffic in a portion of 940 host memory to which the 620 had read-write access via hardware we had built. LaRoy reprogrammed the 620 to send and receive TTY traffic there. With this change the CTE equipment was retired and one cable connected the host to the Tymnet 620 base.

About this time the remote echo issue bubbled to the top of the queue. We also began to introduce forwarding logic in the nodes so that traffic could travel through several nodes, each hop via another leased line. This required routing which was provided by static routing tables. A dial-up access port would reach the same host for weeks at a time, but at least we had much greater flexibility in allocating ports to hosts. Sales spent many painful meetings devoted to this allocation.

We had been thinking about dynamic routing tables to be built as new users logged in. We had no computer with enough power to do this in a brute force way. We could have put it on a Varian machine but that lacked sufficient memory. We chose to program it for the 940 with a slightly modified version of the operating system. LaRoy wrote that program using algorithms to dynamically sense the network topology. That first "Supervisor" knew several versions of the node software and the sizes and absolute addresses of routing tables in each of those versions. The supervisor would send out "data-grams" to cause the distributed nodes to change their routing tables. Several 940s (at least two) would simultaneously host supervisors which would politely agree among themselves which was to supervise the network. If such a supervisor or its host crashed, another supervisor would notice, discover the current net topology and assert control. This might take a few minutes, during which current sessions were unaffected but new users could not log in. Most days went by without such disturbances.

Bringing up the first supervisor.

Tymnet was running well and already serving customers very well. Circuits were laid out manually much as AT&T would manually engineer leased lines to be rented by customers. That was tedious and often involved a few weeks response time. We had vaguely imagined from the beginning of Tymnet the real-time automation of new circuits. The SDS 940 was the largest machine we had available (192 KB) and so we used it to dynamically setup circuits at the beginning of each login session. The new program was called the supervisor.

We converted a small fraction of Tymnet, perhaps 20 nodes around Cupertino, to the new system. With the new system any dialup port would be able to reach any timesharing host. We went thru the normal debugging stages which were eased since the nodes were all local. There were manifest difficulties in arranging yet larger test networks. The flexibility of switched circuits was badly needed. The cost of the switchover included manually reloading a few hundred nodes in many cities. We finally decided in November 1971 to bite the bullet and convert the whole network. One Friday evening as the timesharing load waned, we took down the whole network by beginning to load the new Tymnet node code into each of the many nodes.

The 940 struggled valiantly, reporting the number of nodes whose connections it had discerned. We waited and would take occasional steps to help it along by adjusting heuristic parameters. I am not of much use after midnight and I went home. LaRoy (Tymes) stayed on. Early the next morning I came in and found all of Tymnet up with customers. I assumed that LaRoy, who had written all the code, had found and fixed the problems. An hour or so later LaRoy came in and assumed that I had found the problems that morning. The supervisor had not been able to come up while we were still there and meddling. After LaRoy had finally left, the supervisor did its methodical thing and had mapped and taken over the whole net.

Of course, other bugs were found and improvements were made but Tymnet remained supervised from that point on. One day in a marketing meeting, Virgil Swearingen suggested the name "Tymnet". Few liked the name, but it was clear that it was the best that had yet been suggested.

Development of Tymnet--Part II

As Tymshare acquired more and different host machines, it became even clearer that upgrades were needed to the design of Tymnet. The PDP-10 arrived in 1970 and was adapted to Tymnet by Bill Weiher so as to avoid the expensive "per-user" communications equipment from DEC. Any programmer will tell you that one program knowing addresses in another is a very bad idea, but that was the case in Tymnet 1. It worked pretty well but did limit architectural progress.

LaRoy broached the idea that the supervisor would merely construct a needle which would pass thru the net guiding the creation of a new circuit. The needle would carry a 4-bit nibble, for each node, to steer the needle as it passed that node during the circuit construction. The supervisor would still need to know the network topology and how each node numbered each link. It would also know the node and link loads. It would no longer need to model the routing tables in each node. Routing table entries would be allocated locally. Twenty-five years later this pattern became the main notion behind MPLS routing (Multiprotocol Label Switching). The cost of core memory for nodes had come down and the new node code was larger but more modular and efficient. At this time, nodes typically had 8K or 12K of 16-bit words. A few had 16K.

Once a particularly nasty bug struck the New Orleans node. Usually node bugs would have the civility to strike first, or eventually, near Cupertino, but this one took drastic measures. We loaded an extra core bank with the code for New Orleans and mailed it there in a shoe box. Upon the next crash they swapped core and mailed the core module with the crash state back to Cupertino. We found the bug.

When Tymshare opened the Paris office with its own 940 computer, there was consternation over the idea that the European network would be managed from the United States. This consternation was largely overcome by rewriting a few manuals to describe the operation differently. The Paris 940 ran a supervisor so that when the trans-Atlantic link was out, they would still be in service. Once when all of the American supervisors had gone down, (a very rare occurrence) the Paris node tried to take over the American network via a 2400 bit/sec link. It failed, for by that time 2400 bit/sec did not suffice to control even a night load in the United States.

Tymnet Service

Our timesharing computers attracted business now by those who found it convenient that their program and its data were accessible from a large number of geographical sites. Some customers, however, found our timesharing computers inadequate to their tasks, even with special pricing. There was an increasing call, inside Tymshare and out, to make Tymnet serve hosts other than our own. We gradually did this. This required a considerable programming staff to work with owners of diverse hosts to connect those hosts to Tymnet. Many techniques were used including DMA (Direct Memory Access).

This new business was easy to start because the simplest host interface mechanism was the early cable per connection configuration that required no changes to the host; just power, floor space and termination of a leased line to our site. The prospective customer had very little up-front cost to impede him. After the first day his computer would be accessible by a variety of terminal types from a variety of US cities and even several foreign countries

All through this development, new varieties of terminals were coming to market. They came with higher baud rates and peculiar timing requirements. We would teach Tymnet about these new terminals and the node near the terminal would take care of the peculiarities. Other users of these terminals would have to adapt their host software to accommodate these strange timing requirements. It gradually dawned on us that the adaptation of various hosts to various terminals was a strategic marketing advantage.

With the advent of Tymnet service the Division became a company in its own right in 1979. After McDonnell Douglas bought Tymshare, and then sold it to EDS, Tymnet was sold to British Telecom (who was simultaneously merging with MCI). BT renamed Tymnet as "Global Network Services." At its peak there were 6600 nodes in Tymnet, at nearly as many locations.

Don Johnson, who had been part of Tymnet for many years, said on Monday, March 24, 2003, "We shut down Tymnet at 11:48 AM PST this morning."