

Fermi National Accelerator Laboratory

FERMILAB-TM-2044

The Fermilab Computing Farms in 1997

Stephen Wolbers

*Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510*

February 1998

The Fermilab Computing Farms in 1997

Stephen Wolbers

February 15, 1998

Introduction

The farms in 1997 went through a variety of changes. First, the farms expansion, begun in 1996, was completed. This boosted the computing capacity to something like 20,000 MIPS (where a MIP is a unit defined by running a program, TINY, on the machine and comparing the machine performance to a VAX 11/780). In SpecInt92, it would probably rate close to 40,000. The use of the farms was not all that large. The fixed target experiments were not generally in full production in 1997, but spent time tuning up code. Other users processed on the farms, but tended to come and go and not saturate the resource. Some of the old farms were retired, saving the lab money on maintenance and saving the farms support staff effort.

The year in review

There were some changes made to the farms during the year. As mentioned above, the farms expansion was completed. SGI O200 and IBM RS6000/43P(200 MHz) systems were purchased and integrated into the farms as worker nodes. An SGI O2000 and an IBM RS6000/J50 were purchased as I/O nodes, though only the O2000 was actually integrated into the farms. The J50 was loaned to the mass storage project for use in the HPSS migration. The networking was expanded by replacing the Catalyst 5000 switch with a Catalyst 5500 switch. The O200's, the O2000, and the original Challenge DM were connected to fast Ethernet ports on the switch. The remaining machines were connected to regular Ethernet. The J40 is still connected via FDDI. Additional tapedrives and disk was connected to the I/O nodes as part of the farms expansion.

The collider experiments used very little of the farms in 1997. CDF was essentially completely off the farm by the beginning of 1997. D0 was allocated a piece of the old

farms for most of 1997 but did not make very much use of that allocation after the first few months of the year.

There were many users on the farms in 1997. Many of the fixed target experiments used time to tune up their code and to learn how to use the farms. These include E831, E781, E871 and E872. These experiments were all either close to full production or were in full production at the end of 1997. E866 was in full production at the beginning of 1997 and continued through the year. By the end of the year E866 had completed most of its production.

The other users of the farms came from many parts of the lab and of many areas of physics. The Auger project used the farms to generate substantial shower libraries for their detector studies. Each shower took a substantial amount of CPU time so their use of the fastest processors on the farm (the O200's) was a good match. The Recycler project continued to make use of the DEC portion of the farm when they had a requirement for studies. The Theory group calculations were accommodated. Late in the year a set of calculations for studies of stochastic cooling were started. Other calculations were accomplished on the farms for the MINOS experiment and for magnet calculations.

At the beginning of the year the node allocations still primarily reflected the use of the nodes by the collider experiments and by E706 while it was ramping down. E866 was ramping up and was in full production at this time. By the end of the year the allocation was much different, with the bulk going to E866, E871, E831, E781, E872, the stochastic cooling calculations (Beams), Auger and a little to MINOS and magnet calculations.

The R3000 Indigos were successfully decommissioned near the end of 1997. They were still partially useful, but the memory size, speed, and configuration made them hard to use for larger calculations. Removing them from support saved the lab a large amount of money and saved the support personnel effort in the daily maintenance of those 80 machines. By the end of 1997 the farms consists of about 64 RS6000/320, 34 RS6000/320H, 40 RS6000/220, 70 SGI 4D/35, 45 SGI R5000 Challenge S, 5 SGI O200 (4 processor each), 20 RS6000/43P(133 MHz), 14 RS6000/43P(200 MHz).

CPU utilization

Table 1 provides the summary of CPU time (in VAX-equivalent units) for the whole farm in 1997. A plot of the CPU utilization (including all previous years of the UNIX farms) is shown in Figure 1. The utilization during 1997 was pretty variable

due to the lack of a strong demand over an extended period by any of the users of the farms.

Table 2 and Figure 2 show the utilization for each of the many experiments that have used the farm during 1997. Auger, E866, theory, E831 and E781 were the top users during the year. This reflected the state of their code or the amount of simulation that was required.

Table 3 is a sum of all the CPU time used by all the experiments that have used significant CPU time on the farms during the last 7 years, along with the totals used by each. D0 offline is the largest total user on the farms, with E706 and E791 second and third. CDF is close behind in 4th place. 1998 should see the fixed-target experiments move up in this list as large amounts of computing power are used to process the large datasets written during 1997.

Table 1 – Total CPU use on the Farms – 1997

<u>Month</u>	<u>CPU delivered</u> (Vax-Months/month)
January	2021
February	2011
March	4802
April	3533
May	1848
June	3203
July	2248
August	3897
September	3356
October	2529
November	3896
December	5263

Table 2 – CPU use by experiment – 1997

<u>Experiment</u>	<u>CPU time</u> (Vax-years)
Auger	880
e866	409
Theory	309
e781	302
e831	234
Minos	221
e871	191
recycler	130
Beams	105
d0	96
Magnet	28
e872	21
e706	19
cdf	0
TOTAL	3217

Table 3

Integrated Farm Use
(In units of MIP-years)
Through December, 1997

Experiment	1991	1992	1993	1994	1995	1996	1997	Total
D0 offline		59	570	1072	1184	1743	96	4724
E706	28	82	732	992	795	686	19	3335
E791		100	1232	1249	214	11		2806
CDF		110	320	438	752	956		2576
D0 MC	101	162	396	197	108	23		987
E665	14	105	733	91	2	10		955
E771		94	211	339	219			863
Auger							880	880
E866						55	409	464
E789		156	247					403
E687	99	235	29	30	2			395
Theory							309	309
E781							302	302
E831							234	234
Minos							221	221
Recycler						61	130	191
E871							191	191
Beams							105	105
E760		54						54
E731	38							38
Magnet							28	28
E872							21	21
Total	267	1156	4541	4408	3276	3545	3217	20410




Figure 1.

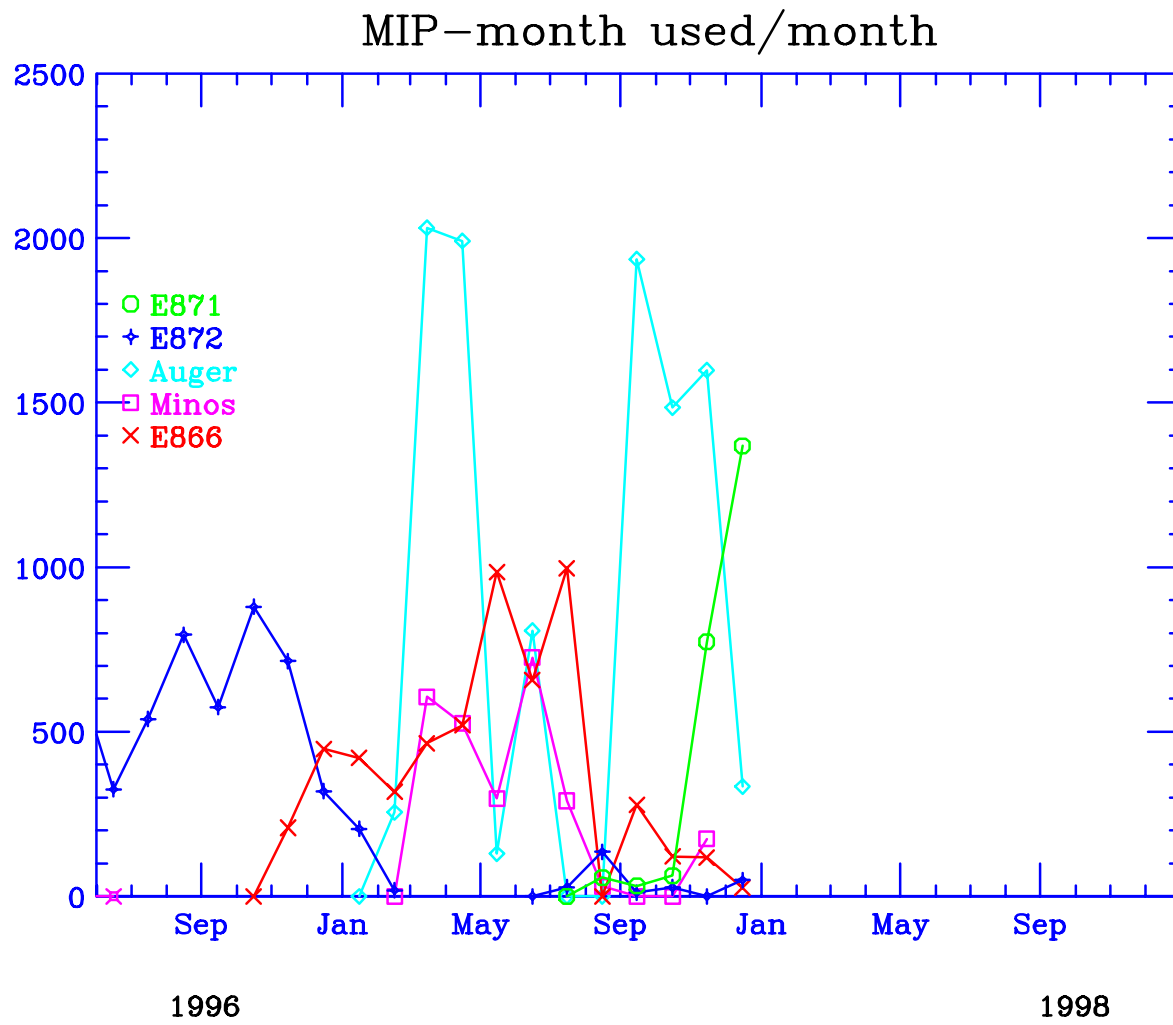


Figure 2(a).




Figure 2(b).

Plans

During 1998 there will be very large demands for CPU time on the farms, old and new. The largest obvious demand for CPU power will be to reconstruct the large datasets that were collected during the fixed target run that ended in September 1997.

The fixed-target experiments will be able to use most if not all of the CPU available on the UNIX farms as they exist today. However, there are other users of substantial CPU time that are either using the farms now or that have need of the CPU time sometime in the near future. Given the existence of such demands the current farms will not be able to deliver all of the CPU that is required.

To accommodate the increased demand there will be an acquisition of a PC farm during 1998. This farm will likely consist of about 20 dual Pentium-II PC's. Linux will be used as the operating system on these systems. The networking will be very similar to that of the current farms (it might even use the same switch). I/O nodes will either be PC's or more traditional UNIX servers. The first use of these PC farms will be fixed-target event reconstruction (to help provide the CPU required to get them finished quickly) and some of the simulations which require large CPU. It is hoped that much will be learned in running a small PC farm.

During 1998 more of the old farms (some of which were purchased in 1991) will be decommissioned. This should not have much of an impact on the total amount of compute power available but it will reduce maintenance costs and effort required to keep the farms functioning.

The next large increase in compute power of the farms will likely occur in 1999 as part of the preparations for Run II. Some of these systems will likely be made available to the fixed-target experiments that plan to take data during 1999. This will allow the fixed-target experiments enough compute power to finish their processing in a reasonable amount of time and will give valuable experience in running a large PC farm.