



**Fermi National Accelerator Laboratory**

**FERMILAB-TM-2006**

## **The Fermilab Farms in 1996**

Stephen Wolbers

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

May 1997

## **Disclaimer**

*This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.*

## **Distribution**

*Approved for public release; further dissemination unlimited.*

## **The Fermilab Farms in 1996**

Stephen Wolbers

January 15, 1997

### **Introduction**

The farms in 1996 began a period of transition. The old farms continue to be used but do not provide sufficient CPU power, memory, or network bandwidth for all of the tasks which are required. Therefore, we have purchased and installed a substantial increment of new farms and are working on adding another increment during 1997. The purpose of all of this activity is to provide computing for the fixed target run and for the other large computing users who cannot be accommodated on the other systems that are available at Fermilab.

### **The year in review**

The old farms were mostly stable during 1996. The largest hardware change was a reconfiguration of the farms required to provide D0 with a 5th I/O server and the worker nodes that were attached to it. In addition, memory was added to most of the D0 worker nodes in order to bring them to 32MB or more.

The farms were largely dedicated to CDF and D0 for the whole of 1996. Other users were still running on the farms, and these users include e706, e791, and e665 from the 1990-91 fixed-target runs and e866 from the 1996-97 fixed-target run. In addition, users from the Accelerator (now Beams) Division running calculations for the Recycler ring and from the Theory Department used the DEC farm during 1996.

CDF and D0 were first able to finish processing the RUN1B and RUN1C data and then work on various additional processing tasks during 1996. CDF was able to process the RUN1C data, data written to tape which was intended to be processed after the end of the run. At the time of this memo the RUN1C data is essentially finished. D0 had many processing tasks, including the D0FIX and WJJ projects. By the end of 1996 most of these were finished.

E791 was able to finish up some Monte Carlo processing (both generation and reconstruction) as well as some splitting. E706 continued to generate and reconstruct Monte Carlo events. E665 finished reconstruction of a set of Monte Carlo events.

E866 was able to begin and make substantial progress on reconstruction of data that was taken in the fall of 1996. The number of nodes allocated to E866 was rapidly ramped up to meet the needs of the experiment.

Two new users were welcomed to the DEC farm. The Recycler project had need for a substantial amount of computing power for calculations related to the design of the Recycler ring. They were allocated the 8 nodes of the DEC farm, a farm which was purchased in FY95 and was commissioned during 1996. At the end of the year a need for Theory calculations was satisfied by using the DEC farm, idle until early 1997 when the Recycler calculations will begin again.

The node allocations to various experiments at the beginning of the year were : CDF-97, D0-116, E706-45, E791-28, E665-4, E687-5. At the end of the year the allocations of the old farm nodes was: CDF-97, D0-122, E706-16, E866-41. The DEC farm (8 nodes) was allocated to the Recycler Project.

There was a great deal of progress on the software used on the farms. cps was upgraded and a new version (3.0) was built, tested and released in 1996. This new version is able to handle 64-bit operating systems and has an interface to the tape operator system OCS. A new cps\_batch was designed and written. The new cps\_batch retains most of the functionality of the old cps\_batch and even carries forward most of the commands. The underlying database is new and the operator interface is xtape rather than the old xoper screen. The transition to the new cps\_batch is almost complete. These are major software accomplishments.

Finally, the oldest UNIX farms, the original 25 4D/25 nodes, were retired on February 7, 1996. The rest of the farms were moved on the second floor of Feynman on April 9-10, 1996. This was a big move and the ability to bring the systems back up quickly after the move was due in large part to the excellent planning done by Randy Hilgers, Phil Lutz, Khalid Paden, Chuck Andrews, Steve Fry, Reggie Gibbons, Tom Naughton, Mark Kaletka, Marc Haibeck, Lisa Giacchetti and Jack MacNerland.

## **CPU utilization**

Table 1 provides the summary of CPU time (in VAX-equivalent units) for the whole farm in 1996. A plot of the CPU utilization (including some previous years) is shown in Figure 1. The utilization was fairly steady and followed from the fact that both CDF and D0 were processing large datasets which were already on tape. The

fixed-target experiments were using CPU whenever it was available (E706) or when the data became available on tape (E866).

Table 2 and Figure 2 show the utilization for each of the many experiments that have used the farm during 1995. D0 offline was the largest user. E706 and CDF were both large. All other users were either finishing up and used very little CPU time or were just getting started at the end of the year.

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 5 years, along with the totals used by each. D0 offline is now the largest total user on the farms, with E706 and E791 second and third. CDF is close behind in 4th place. It is interesting to note that both CDF and D0 have used a larger amount of computing time on the farms during any given year than the previous year. This trend should end this year as both experiments are essentially finished with farm processing.

Table 1 – Total CPU use on the Farms – 1996

<u>Month</u>	<u>CPU delivered</u> (Vax-Months/month)
January	3324
February	3709
March	2957
April	2788
May	3406
June	3976
July	3559
August	4725
September	5076
October	4859
November	3402
December	1742

Table 2 – CPU use by experiment – 1996

<u>Experiment</u>	<u>CPU time</u> (Vax-years)
d0(total)	1766
d0(Reconstruction)	1743
d0(Monte Carlo)	23
cdf	956
e706	686
e771	219
recycler	61
e866	55
e791	11
e665	10
TOTAL	3545

Table 3

**Integrated Farm Use**  
(In units of MIP-years)  
Through December, 1996

Experiment	1991	1992	1993	1994	1995	1996	Total
D0 offline	0	59	570	1072	1184	1743	4628
E706	28	82	732	992	795	686	3316
E791	0	100	1232	1249	214	11	2806
CDF	0	110	320	438	752	956	2576
D0 MC	101	162	396	197	108	23	987
E665	14	105	733	91	2	10	955
E771	0	94	211	339	219		863
E789	0	156	247				403
E687	99	235	29	30	2		395
Recycler						61	61
E866						55	55
E760	0	54					54
E731	38	0					38
Total	267	1156	4541	4408	3276	3545	17193

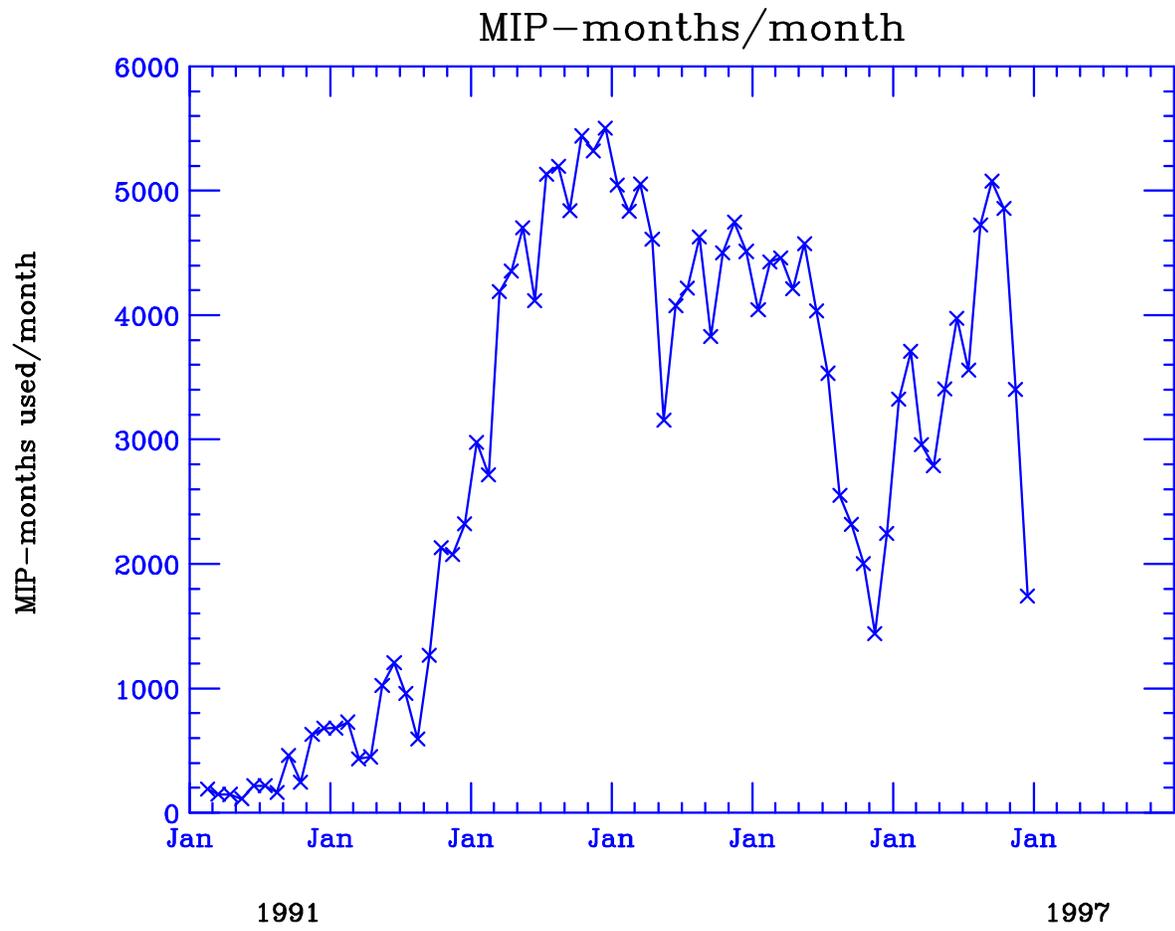


Figure 1.

# MIP-month used/month

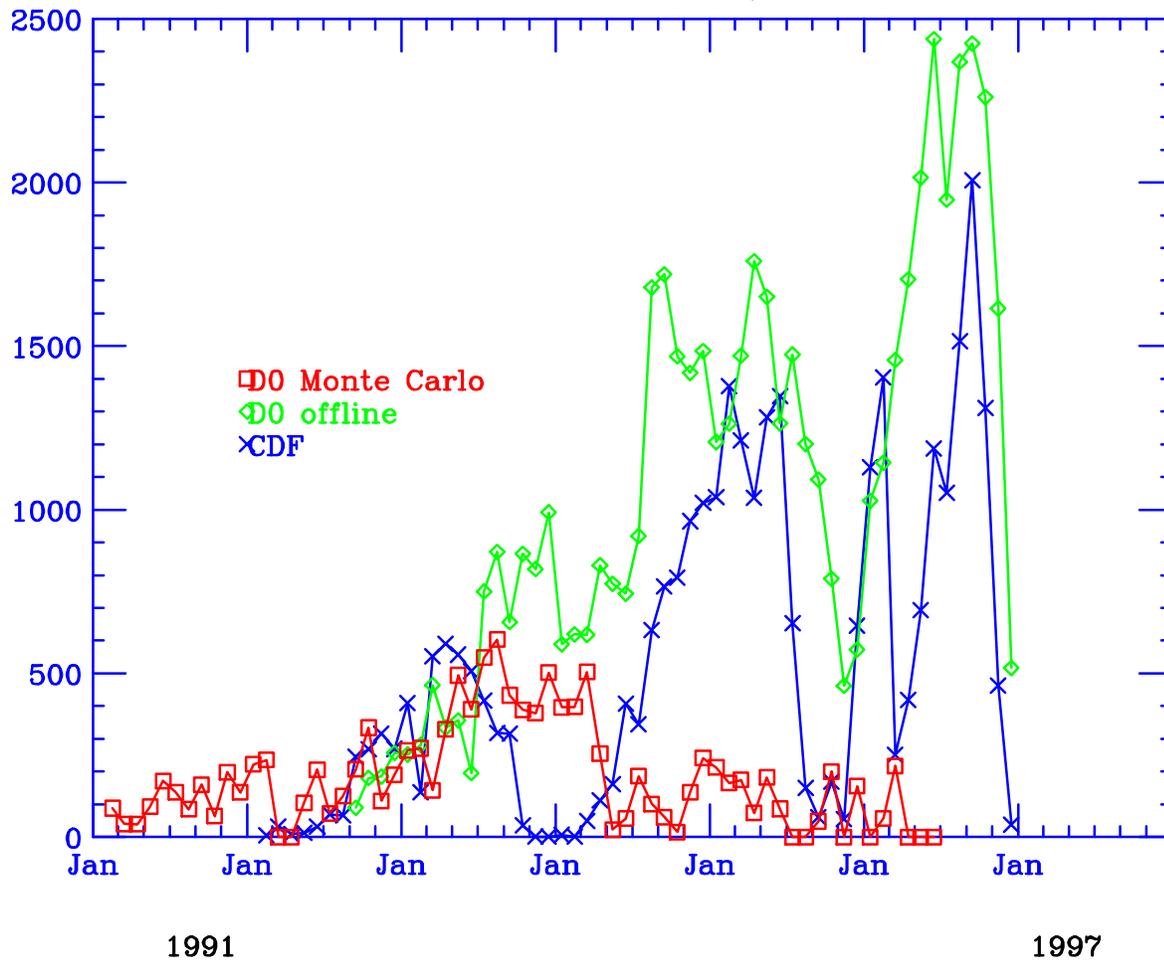


Figure 2(a).

# MIP-month used/month

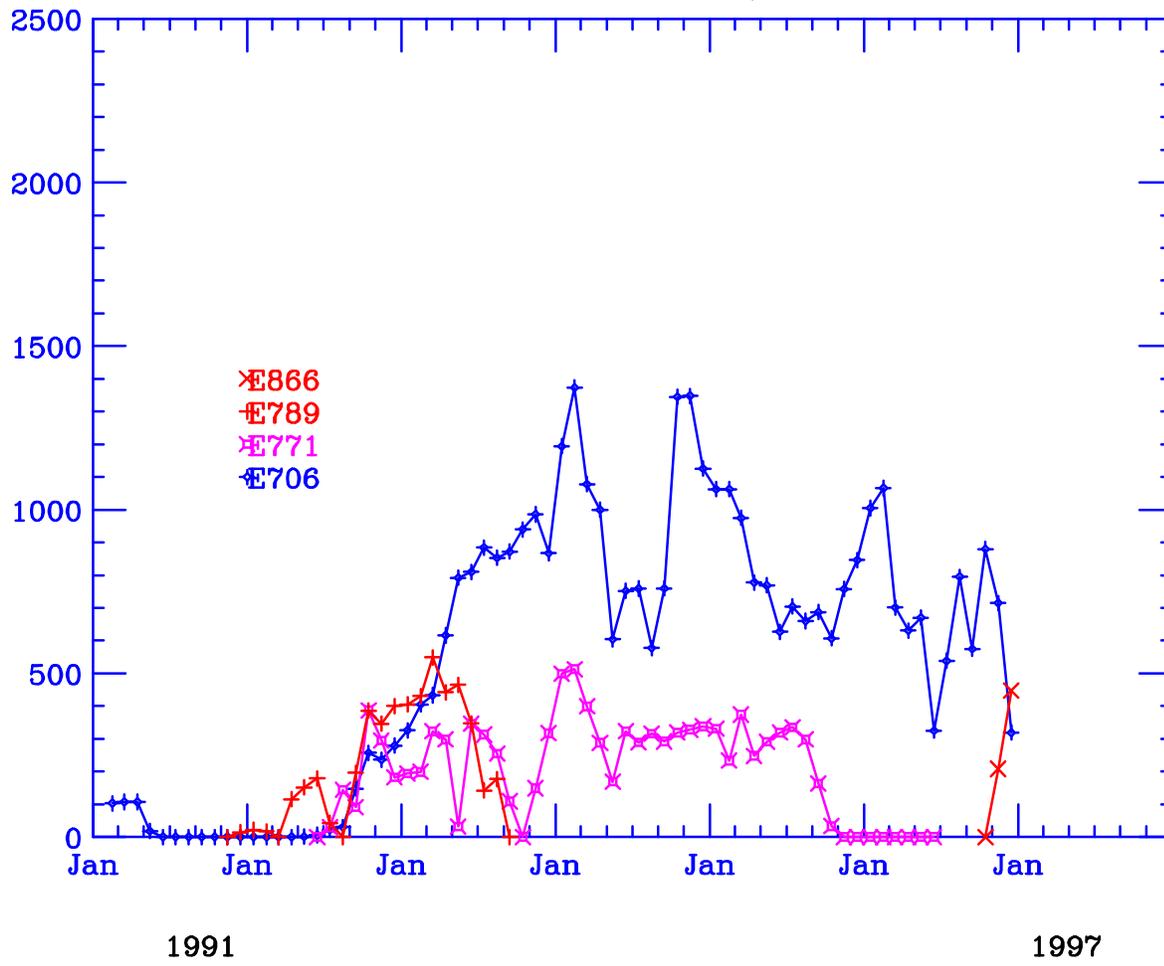


Figure 2(b).



## **The Farm Expansion**

During 1996 a major expansion of the farms was initiated. The goal of the expansion was to provide 7,500 MIPS in FY96. A committee was appointed by Joel Butler to specify and procure the necessary computing. An RFP was written and bids were secured from DEC, IBM, SGI and SUN. SGI and IBM were the first and second place winners and so were awarded 5000 and 2500 MIPS, respectively. The machines were all delivered by early October and the installation and commissioning has been proceeding steadily since then. All the nodes are expected to be installed and running very soon.

## Plans

During 1997 there will be very large demands for CPU time on the farms, old and new. The fixed-target experiments are taking high quality and large quantities of data. It is expected that many of the experiments will be using the farms to process their data during 1997 and into 1998. This will be the highest priority on the farms. Other groups will make use of the CPU time available on the farm, with lower priority.

CDF and D0 will finish whatever processing they still have. This is expected to be small. The resources now used by CDF and D0 will essentially be freed up and will be reprogrammed to other users.

Substantial simulation demands will be accommodated on the farms. The Recycler Project is expected to use the DEC farms for at least a year to perform beams simulation for the Recycler ring. The Theory Department has some very substantial calculations which will be done on the farms. CMS, Minos and the Auger Project all have demands for large amounts of CPU which might best be accommodated on the farms. We hope to make efficient use of the CPU to accomplish all of the tasks.

A second farm buy will proceed during 1997 to provide another 7,500 MIPS of processing power. The purchase will be made from SGI and IBM in the same 2:1 ratio as the initial bid. It is expected that the newest nodes will be available for use by spring of 1997.

It is expected that 1997 will be a transition year for the farms. It will be a transition from CDF/D0 to fixed-target and a transition from the old farms to the new. Our goal is to provide large amounts of computing for those users with large requirements and to provide it as easily and conveniently as possible, keeping in mind the priorities of the lab as well as the desire to utilize the computing as efficiently as possible.