CONF-770525--3

LA-UR-77-1069

TITLE: THE CRAY-1 AT LASL: AN UPDATE

AUTHOR(S): W. JACK WORLTON

SUBMITTED TO: AESOP XVI CONFERENCE
MAY 2 - 5, 1977

los alamos
scientific laboratory
of the University of California
LOS ALAMOS, NEW MEXICO 87544

An Affirmative Action/Equal Opportunity Employer

Form No. 836
St. No. 2629
1/75

UNITED STATES
ENERGY RESEARCH AND
DEVELOPMENT ADMINISTRATION
CONTRACT W-7405-ENG. 36

THE CRAY-1 AT LASL: AN UPDATE

J. Worlton*

Los Alamos Scientific Laboratory

1. INTRODUCTION

The first CRAY-1 computer was installed at Los Alamos Scientific Laboratory (LASL) in March 1976; it was evaluated from April 1 to September 30, 1976, under the guidance of ERDA Headquarters and with the cooperation of the Federal Computer Performance Evaluation and Simulation Center (FEDSIM). Since October 1976 LASL has been integrating the CRAY-1 into the LASL Integrated Computer Network (ICN), developing software, and converting application codes for early productive use of the CRAY-1. This report reviews the background and current status of this computer.

2. VECTOR PROCESSORS: DEFINITIONS AND PERSPECTIVE

The CRAY-1 combines both scalar and vector architecture features in its design. A "scalar" operation is one that typically uses two operands, executes a single operation, and produces one result. A "vector" operation is one that typically uses two sets of operands, executes the same operation pairwise on the elements of these sets, and generates a set of results. An important limitation on the execution bandwidth of scalar operations is the need to issue at least one and generally several instructions for each result generated; this limitation is removed in vector operations, since a single vector command may generate many results.

Data flow in vector architectures has been implemented in two ways: memory-to-memory (MM) and register-to-register (RR). The first generation of vector computers (the CDC STAR-100 and the Texas Instruments ASC) were MM designs in which the vector operands were fetched from memory and the results returned to memory. The CRAY-1 is an RR vector design in which operands are fetched from registers and the results returned to registers. Whereas the MM designs typically have a rather long startup time for the vector operations (several microseconds), the RR design reduces startup time significantly and therefore allows efficient processing of short vectors. Further, with operands and results stored in registers, parallel execution of vector operations can readily be implemented, thereby increasing vector execution rates still further.

---

*Work performed under the auspices of the U.S. Energy Research and Development Administration.

In spite of the fact that vector processing is usually many times faster than scalar processing in any given scalar-vector design, the speed of the scalar processing unit remains critical to the overall performance of such a machine [1]. The reason for this is that very few problems can be vectorized totally, and even a small percent of the results being generated in a slow scalar mode can degrade the performance of the machine to unacceptable levels. Even an infinitely fast vector processor in combination with a slow scalar processor would give inadequate performance unless the applications were almost totally vectorized--and problems of this type are uncommon. And even if the application itself is highly vectorizable, the system software which performs service functions for the application code is typically implemented in scalar code.

3. THE CRAY-1: SUMMARY OF CHARACTERISTICS

The CRAY-1 is a logical descendant of the CDC 7600, and Table 1 summarizes some basic characteristics of these two computers. In some respects the CRAY-1 design simply improved on the 7600, e.g., a faster minor cycle time, improved instruction buffering, additional scalar registers, and faster memory cycle. Other features of the CRAY-1 are more accurately characterized as "breakthroughs", e.g., the cooling technology, the inclusion of the RR vector mode and vector chaining, and the two-to-three orders of magnitude increase in the capacity of bipolar memory. These items are radically different in either concept or quantity than those found in prior computers.

Parallelism is implemented in the CRAY-1 in several ways: in the pipelined arithmetic units, each of which can have several operations in process at one time; in the inclusion of multiple arithmetic units which can operate in parallel; in the chaining feature which allows results to be used as operands essentially as soon as they are generated without waiting for the vector commands to be completed; in the overlap of instruction processing with functional unit operations; and in the multiple independent data transmissions of the 12 full duplex data channels. In spite of this high degree of parallelism, the machine can be programmed in serial mode without regard to the parallel operations of the various units.

Further details of the CRAY-1 design can be found in [2].

4. EVALUATION SUMMARY

The six-months evaluation of the CRAY-1 at LASL was deemed necessary by ERDA in order to assure that this computer was a sufficient technological advance over its predecessors to warrant acquisition by LASL and other ERDA laboratories. Table 2 lists the threshhold criteria used in this evaluation. The scalar criterion was established in order to avoid the degradation of overall performance that occurred in some prior computers through the combination of slow scalar performance with fast vector performance, thereby severely limiting the class of problems to which such computers could be applied. The vector criteria were

	CDC 7600	CRAY-1
1. CENTRAL PROCESSING UNIT (CPU)		
• MINOR CYCLE TIME (NS)	27.5	12.5
• SCALAR PROCESSING RATE	1	2 - 4
• VECTOR PROCESSING RATE	1	2 - 15
• INSTRUCTION BUFFER	1 X 12 WORDS	4 X 16 WORDS
• EXPLICIT VECTOR COMMANDS	NO	YES
• AUTOMATIC NORMALIZATION	NO	YES
• VECTOR REGISTERS	NONE	8 X 64 WORDS
• SCALAR REGISTERS	24	144
2. MEMORY		
• WORD LENGTH (BITS)	60	64
• CAPACITY (K = 1024 WORDS)	64K + 512K	1024K
• CYCLE TIME (NS)	275/1600	50
• ACCESS TIME (NS)	220	125
• TECHNOLOGY	CORE	BIPOLAR

TABLE 1. COMPARISON OF THE CDC 7600 AND THE CRAY-1

SCALAR OPERATIONS	≥ 2 X CDC 7600 IN SPEED
VECTOR OPERATIONS	
• SHORT (L = 20)	≥ 3 X CDC 7600 IN SPEED
• MEDIUM (L = 100)	≥ 4 X CDC 7600 IN SPEED
• LONG (L = 500)	≥ 5 X CDC 7600 IN SPEED
RELIABILITY	
• MTTF	≥ 4 HOURS
• MTTR	≤ 1 HOUR
• AVAILABILITY	≥ 0.8

TABLE 2. QUALIFICATION CRITERIA

defined for several vector lengths that would allow evaluation of vector performance on relatively short vectors. Three reliability criteria were used instead of the single "availability" criterion typical of GSA schedules.

The details of the methodology and the results of the evaluation can be found in [3], and we list here only the main results of the evaluation. Figure 1 shows the relative scalar performance of the CRAY-1 to the CDC 7600 for 26 kernels coded in assembly language. The relative speed ratios ranged from 2.03 to 4.12. The low end of this range comes from 7600 codes that were contained in Small Core Memory (SCM) and hence were highly efficient on the 7600; the high end of the range comes from codes which required the use of the 7600 Large Core Memory (LCM).

The vector execution rates of the CRAY-1 and the 7600 were compared by programming five vector operations for both machines and comparing the times, with the averages of these timings for the specified vector lengths being compared to the vector qualification criteria. Table 3 shows these results presented in millions of floating-point operations per second (MFLOPS)[a]. The average speed ratios for the CRAY-1 compared to the 7600 were 3.39, 4.50, and 5.12 for vector lengths of 20, 100, and 500, respectively. Thus, the qualification criteria were met in all instances. However, this was not an attempt to determine the maximum or even the expected speed ratios--only an attempt to determine a threshhold level of performance in vector mode.

LASL work with linear system solvers has shown that in favorable programming environments, the CRAY-1 can produce 64-bit results in the 100 to 140 MFLOPS range [4].

The mean-time-to-failure (MTTF) of the CRAY-1 was evaluated by running an "EXERCISER" code designed to access as many different hardware units of the machine as feasible for runs of many hours each, and checking the MTTF over a period of twenty consecutive working days. Mean-time-to-repair (MTTR) and availability were defined in the usual manner. In view of the lack of an adequately demanding workload on the CRAY-1 in this period, the use of the EXERCISER was deemed necessary in order to impose demands as great or greater than those found in a production environment.

---

[a]. The notation for these operations is:

VVSPVS means (vector) = (vector)*(scalar) + (vector)*(scalar)
VVVPVV means (vector) = (vector)*(vector) + (vector)*(vector)
VVVPV means (vector) = (vector)*(vector) + (vector)
DOTPRO means dot product
ADDVEC means (vector) = (vector)*(scalar) + (vector)

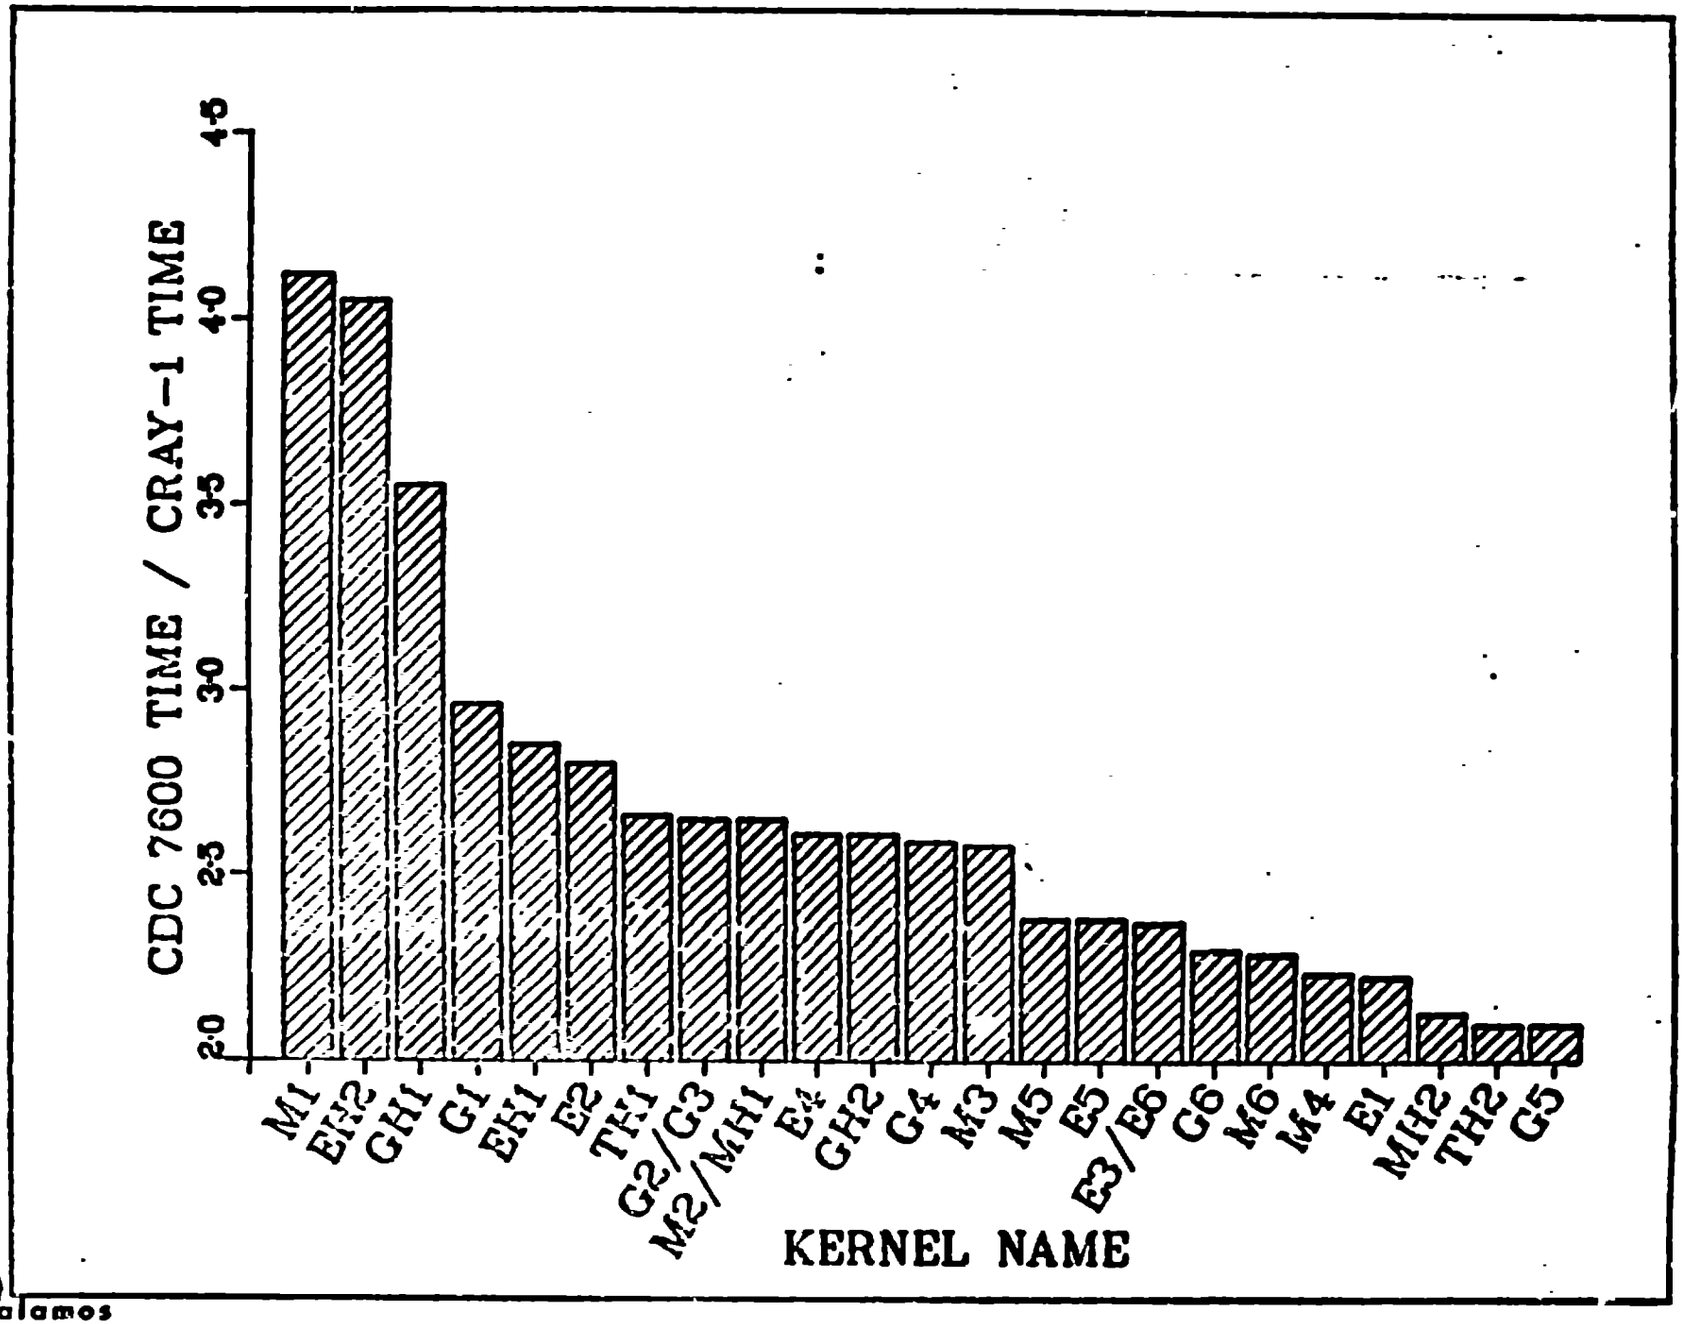

los alamos scientific laboratory

FIGURE 1. SORTED HISTOGRAM OF KERNEL SPEED RATIOS

	L=20		L=100		L=500	
FUNCTION	CRAY-1	7600	CRAY-1	7600	CRAY-1	7600
VVSPVS	31.6	8.2	55.9	11.1	65.6	11.9
VVVPVV	25.5	7.0	38.1	8.5	41.6	9.0
VVVPV	18.8	5.5	30.0	7.1	34.0	7.5
DOTPRO	12.8	6.2	36.3	9.1	60.6	10.1
ADDVEC	22.2	5.6	38.2	8.0	43.9	8.9
AVERAGE	22.2	6.5	39.7	8.8	49.1	9.5

TABLE 3. CRAY-1 AND 7600 PERFORMANCE IN MFLOPS

Figure 2 shows the MTTF data for successive twenty-day periods. All three reliability criteria were required to be satisfied for a period of twenty consecutive working days, and this happened for many periods. For example, in the period 5-14-76 to 6-11-76, the MTTF was 7.08 hours, the MTTR was 0.38 hours, and the system availability was 95%.

Most of the failures detected during the evaluation period were transient memory parity errors; of the total of 170 errors, 152 were memory parity errors, and of these, 140 were transient (i.e., non-reproducible). Since the evaluation period has been completed, Cray Research has offered this computer with single-bit error correction and double-bit error detection (SECDED) as an option; if this feature had been in operation during the evaluation period, it would have extended the MTTF significantly. For example, if SECDED had eliminated the single-bit transient errors, the MTTF would have been increased by a factor of 5.5; if it had also eliminated the reproducible single-bit errors during production runs and made it possible to correct these errors during preventive maintenance periods, the MTTF would have been increased by a factor of 9.4.

5. CURRENT STATUS AND PLANS

5.1 Hardware

Configuration. The CRAY-1 configuration installed at LASL for the evaluation period included the central processing unit with 512K (K=1024 64-bit words) of memory, one disk control unit, two DD-19 disks, and a Data General Eclipse that was used for maintenance and front-end functions. This configuration has recently been upgraded to three disk control units and five DD-19s. It is planned to upgrade the main memory capacity to 1024K words including the SECDED feature during third quarter 1977; further disk control units and disks will be added to provide an eventual configuration of nine controllers and seventeen disks.

Networking. The CRAY-1 has been connected to one of the LASL 7600s running the LTSS operating system in order to provide access to the network resources and to allow LTSS to be used as a support processor for the CRAY-1. In the long term, other computers in the network, including those of smaller scale, will be able to use the CRAY-1 as a resource for executing large-scale tasks, with the pre- and post-processing for these tasks being done on the smaller-scale machines. In this way, it will be possible to minimize software development work and to bring the CRAY-1 up to productive status much earlier than would be possible if a full software system were developed on the CRAY-1 itself.

Reliability. The EXERCISER code has continued in use to serve as a check on the trend in CRAY-1 reliability. However, changes in the system environment have made it necessary to reduce the size of the linear system being solved from 705x705 to 561x561 in order to fit the code and data into the user area of memory. Thus, this form of the code (EXERCISER-II) checks less of the

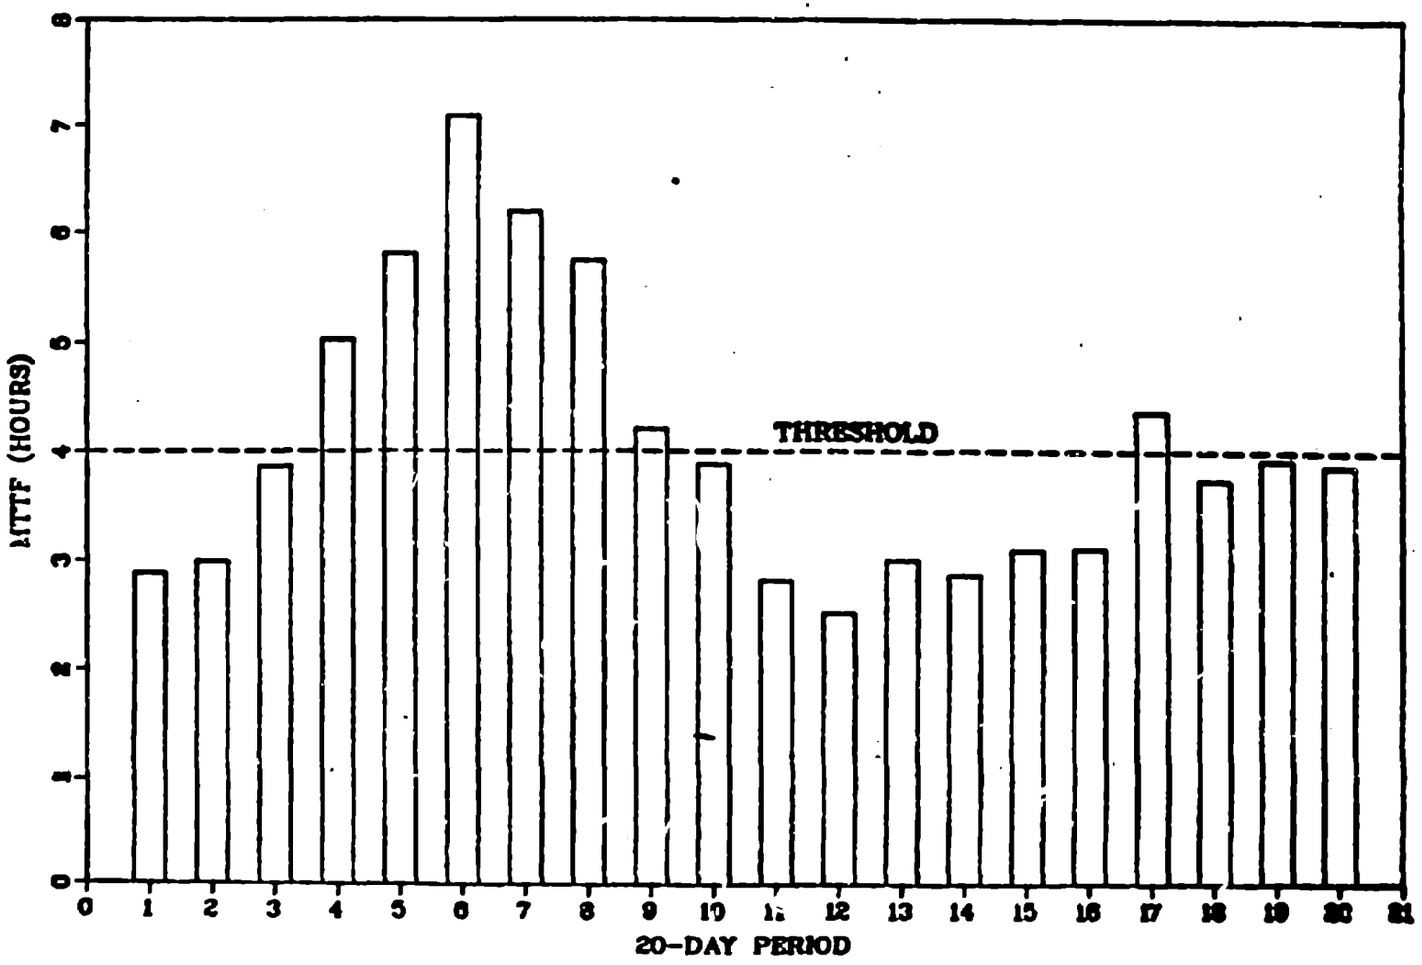

FIGURE 2. MTTF FOR CRAY-1 EVALUATION

memory than its predecessor and the data is not directly comparable. However, there is a significant trend upward in the MTTF data for EXERCISER-II, as shown in Table 3.

MONTH	MTTF (HOURS)
January	5.16
February	10.31
March	18.22
April (to 24th)	53.13

TABLE 3. EXERCISER-II Data for 1977

In addition to the growing experience of the CRAY-1 maintenance engineers, the primary cause of this improvement is the use of wider voltage margins during preventive maintenance periods to locate and remove marginal components.

5.2 Software.

5.2.1 Software on LTSS to support the CRAY-1. As noted above, the LTSS system is being used to support the CRAY-1, including the following software products.

CRAYFTN. This is a FORTRAN cross-compiler for the CRAY-1; it is a LASL-developed compiler based on the CDC FTN 4.6 compiler. It is run only on FTN-licensed machines and it is maintained solely by LASL. This compiler is also available under the NOS system. The output of this compiler is assembly code for the Cray Assembly Language (CAL). Work is continuing on the optimization phase of this compiler.

MODEL. MODEL is a systems-implementation language for the CRAY-1 developed at LASL. MODEL is also a cross-compiler.

CRALINK/CRAYCNTRL. This is software to control the link between the CRAY-1 and LTSS. These routines accept CRAY-1 jobs from the LTSS users, queue up jobs, execute these jobs serially on the CRAY-1, and return CRAY-1 output to the user.

CRACNV. This routine converts data between LTSS and CRAY-1 file formats.

CRAEDIT. This is a utility program to allow LTSS users to examine CRAY-1 format files on LTSS without converting files to LTSS format.

5.2.2 CRAY-1 software. On the CRAY-1 itself, the following software products are available.

Loader. This is a relocatable loader with overlay capability that can satisfy externals from multiple libraries.

CAL Assembler. This is a macro assembler with features similar to the CDC COMPASS assembler.

BEDIT. This is a utility to manipulate binary programs and subprograms. It allows users to maintain a library of binary programs and subprograms, thus eliminating the requirement to recompile a whole production code.

DISPOSE. This is a utility to route CRAY-1 files to LTSS.

DMPDS. This utility dumps the contents of a CRAY-1 data set in octal or other formats.

DMP. This is a utility used to dump the contents of CRAY-1 memory in octal or other formats.

LIBMAK. This is a utility to construct libraries for the loader.

UPDATE. This utility is used to maintain and modify source programs.

5.2.3 Operating system status and plans. The operating system currently in use on the LASL CRAY-1 is a modified version of the operating system developed by Cray Research. In the long term, this system will evolve into a LASL-developed operating system (DEMOS) that will provide a hierarchical file system, multitasking (multiprogramming), process structure, a flexible job control language, and a network interface. The network interface will continue the philosophy of the "support processor," whereby other smaller computers will be able to use the CRAY-1 as a "computational engine."

5.3 Code conversion.

The LASL philosophy of code conversion for the CRAY-1 is to convert a severely-limited set of high-priority production codes initially, and to develop the system resources needed by these codes. With these codes in production status, the system resources can then be expanded to meet the needs of a larger set of codes.

Consistent with the justification for the acquisition of the CRAY-1 at LASL, the initial set of codes has been taken from the

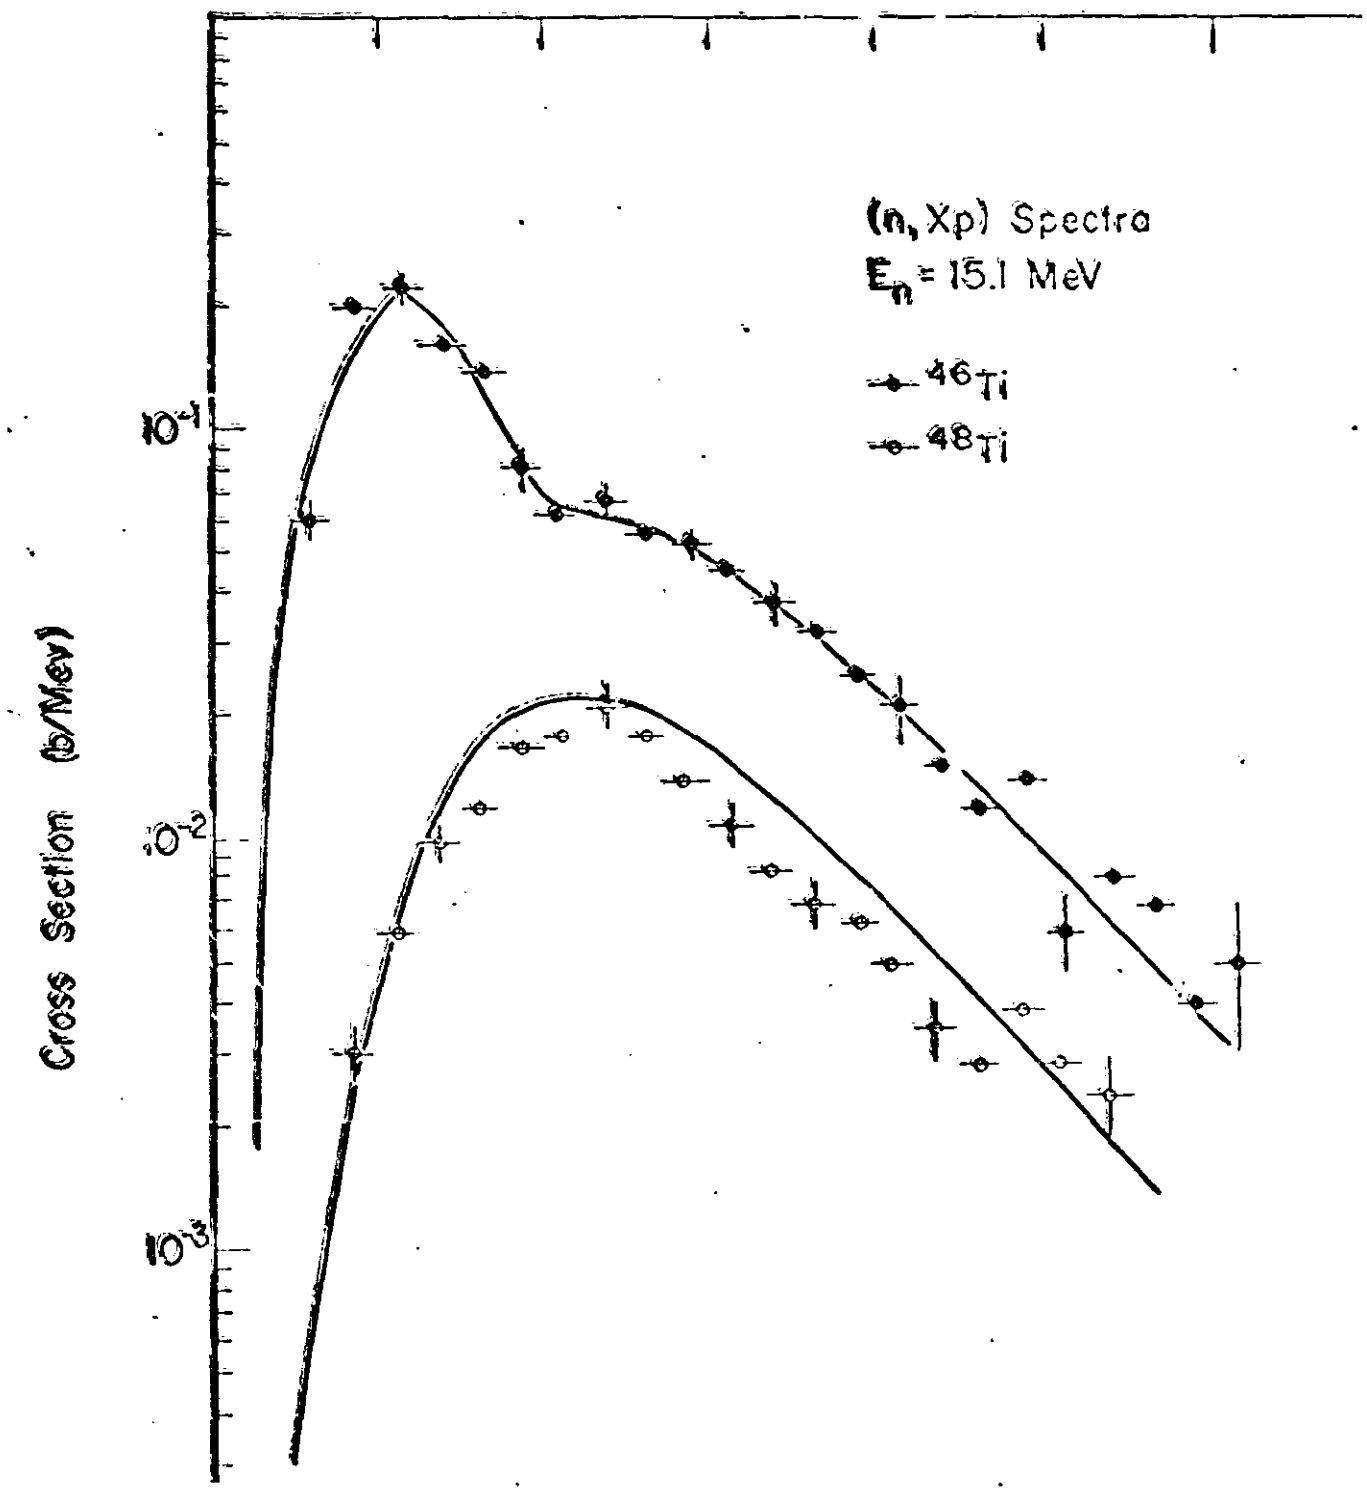
(n, Xp) Spectra
E_n = 15.1 MeV
^{46}Ti
^{48}Ti
Cross Section (b/Mev)
10^{-1}
10^{-2}
10^{-3}

weapons program, and in this set the code known as BIGMAC has been used as the primary definition of system development needs. A working subset of this code has been running correctly since late in 1976; numerous full-size problems have been run. Timing studies of BIGMAC have shown that the CRAY-1 scalar version of the code runs 3.1 times as fast as the same code running on the 7600 with the RUN compiler under the CROS operating system. Vectorization of this code is proceeding, in order to further speed up the execution on the CRAY-1.

Five other weapons codes are in various stages of conversion, with several of them largely converted. The Monte-Carlo code MCN is running correctly on the CRAY-1, tracking the LTSS and CROS versions of the code.

In addition to the conversion of these weapons codes, conversion of codes in four other areas is in process. The LASL Testing Division is converting a two-dimensional hydrodynamics and transport code used for underground containment studies; a three-dimensional version of the code is planned for the CRAY-1. A three-dimensional magnetohydrodynamics plasma simulation code is being converted for use in magnetic fusion energy studies. A reactor safety code (SIMMER) is being converted for work being done for the Nuclear Regulatory Commission (NRC). Finally, a research code for statistical mechanics has been converted. As noted above, these non-weapons codes are being prepared and run on a non-interference basis with respect to the work on the weapons codes.

6. SUMMARY

The CRAY-1 has been evaluated under carefully controlled conditions and has been shown to have performance and reliability characteristics that are well matched to the programmatic needs at LASL. Software development, networking, and code conversion have proceeded to where the CRAY-1 is now in early production status. With the upgrading of the memory to one million words, the upgrading of the disk configuration, and the continuing development of system and applications software, this computer will become the most powerful worker computer in the LASL Integrated Computer Network.

REFERENCES

1. Rudsinski, L., and Worlton, J., "The Impact of Scalar Performance on Vector and Parallel Processors," Symposium on High-Speed Computer and Algorithm Organization, Academic Press, April 1977.

2. The CRAY-1 Computer System, Reference Manual 2240004, Cray Research, Inc., May 1976.

3. Keller, T. W., "CRAY-1 Evaluation Final Report," Los Alamos Scientific Laboratory, Los Alamos, N. M., LA-6456-MS, December 1976.

4. Jordan, T. L., and Fong, Kirby, "Some Linear Algebraic Algorithms and Their Performance on the CRAY-1," Symposium on High-Speed Computer and Algorithm Organization, Academic Press, April 1977.

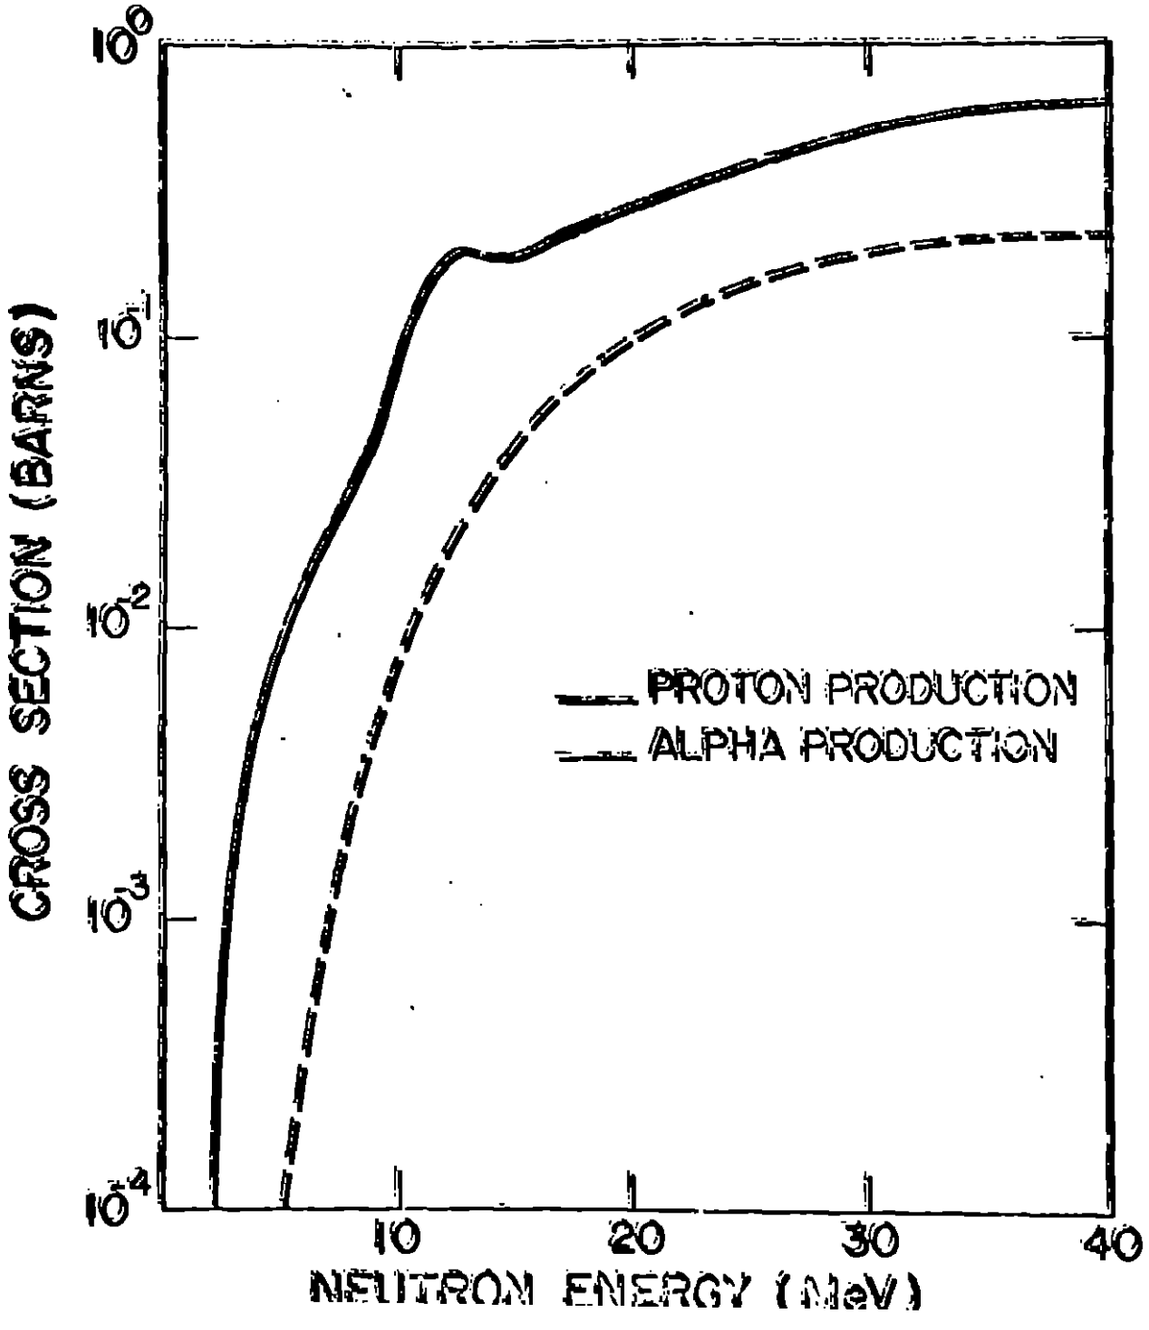
CROSS SECTION (BARNS)
10^0
10^{-1}
10^{-2}
10^{-3}
10^{-4}
PROTON PRODUCTION
ALPHA PRODUCTION
10
20
30
40
NEUTRON ENERGY (MeV)

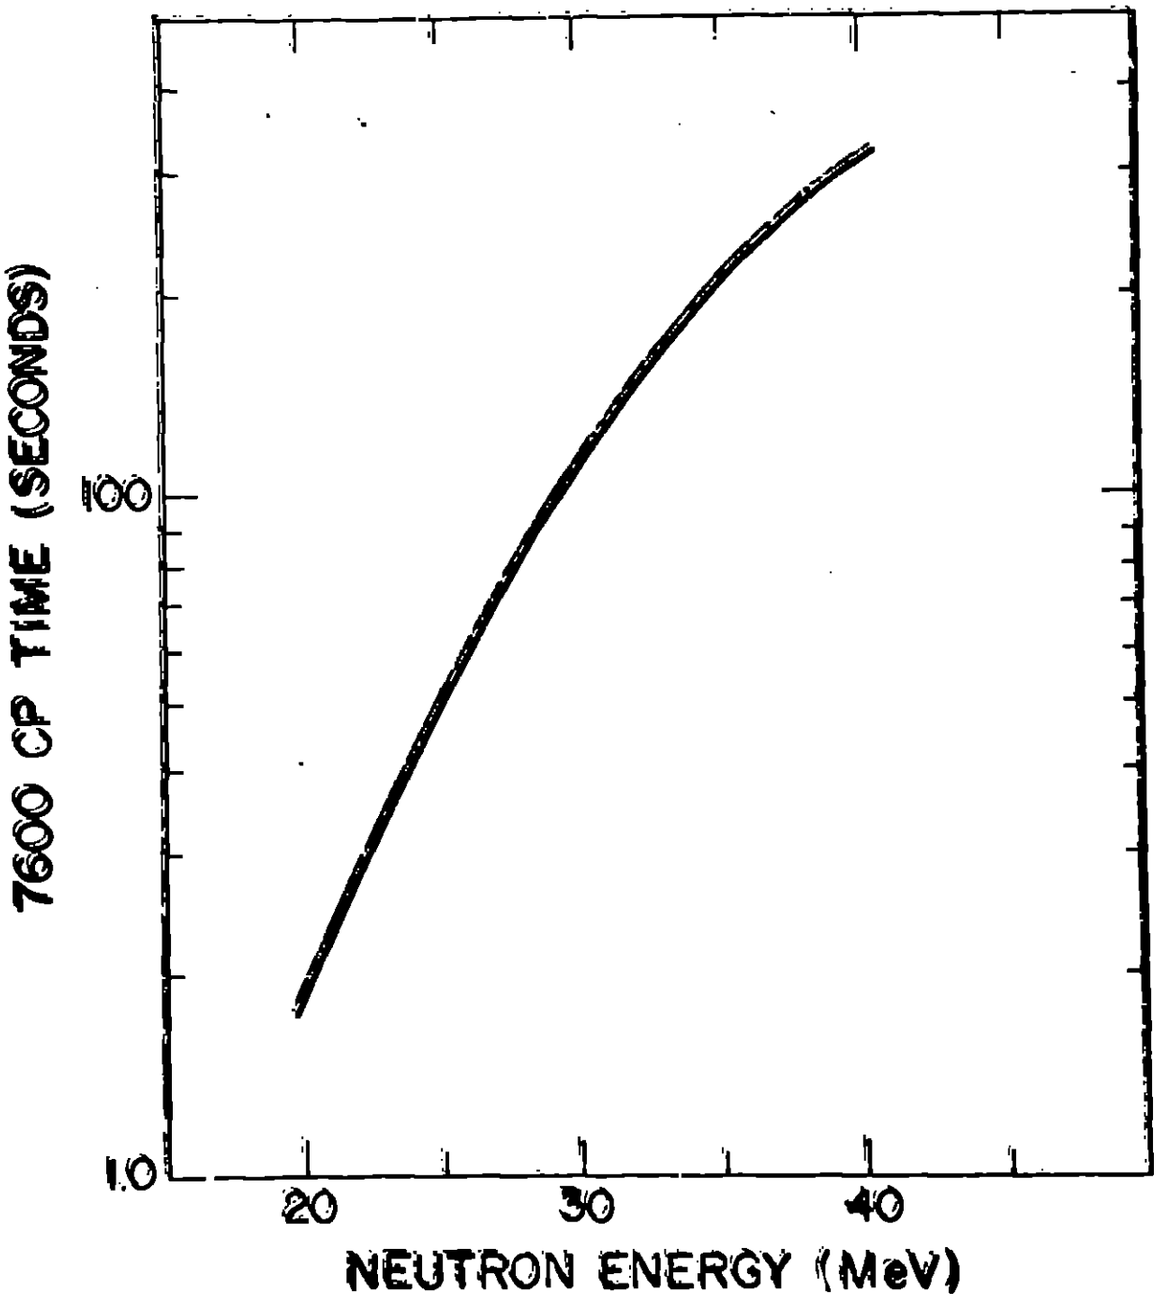
7600 CP TIME (SECONDS)
100
1.0
20
30
40
NEUTRON ENERGY (MeV)

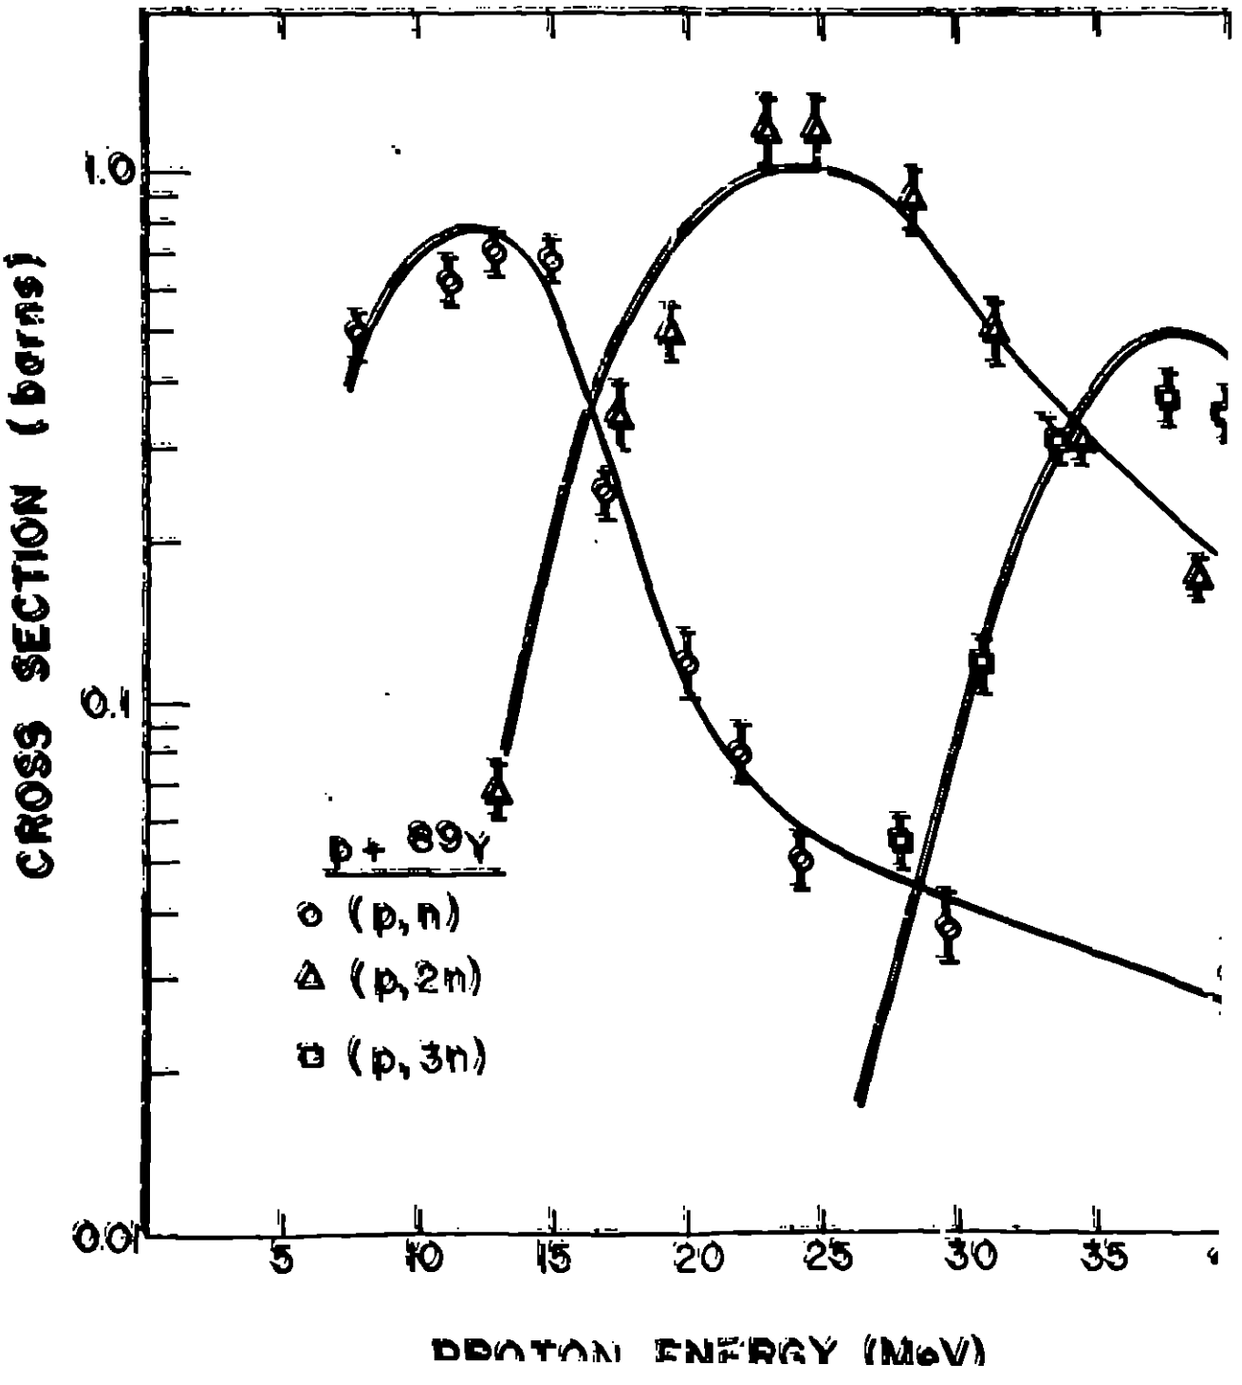
CROSS SECTION (barns)
1.0
0.1
0.01
p + 89Y
o (p, n)
Δ (p, 2n)
□ (p, 3n)
5
10
15
20
25
30
35
PROTON ENERGY (MeV)

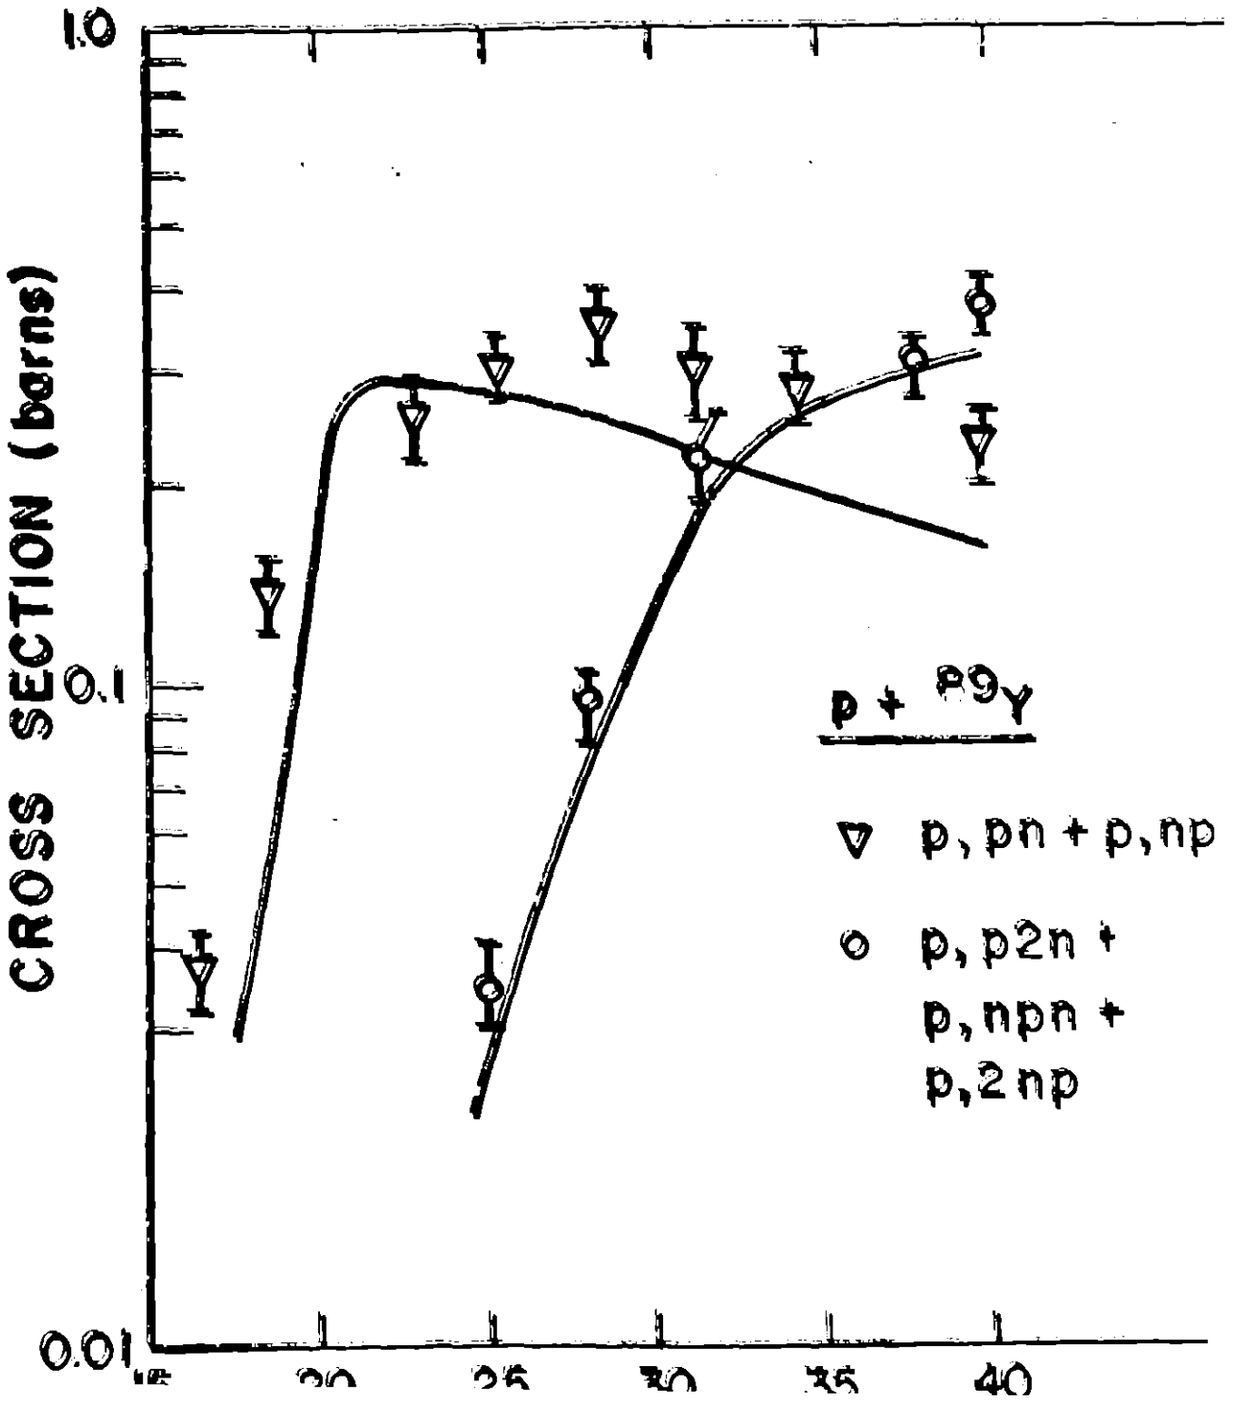
1.0
0.1
0.01
CROSS SECTION (barns)
p + 89Y
p, pn + p, np
p, p2n + p, npn + p, 2np
35
40

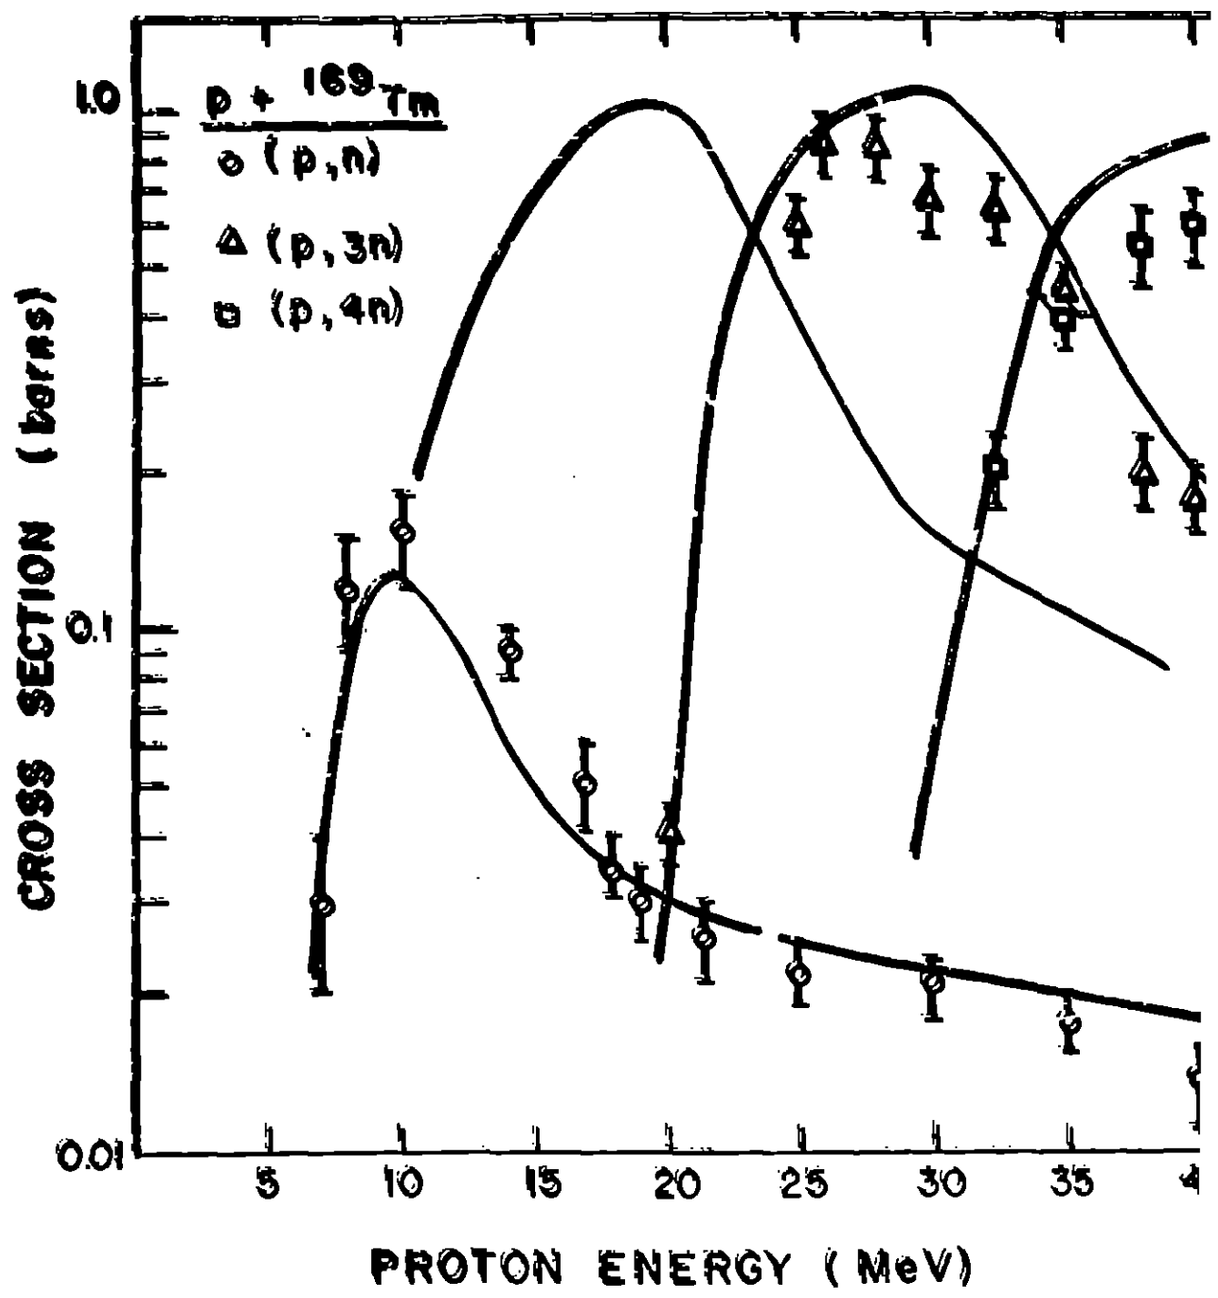
p + 169Tm
o (p,n)
Δ (p,3n)
□ (p,4n)
CROSS SECTION (barns)
1.0
0.1
0.01
5
10
15
20
25
30
35
PROTON ENERGY (MeV)

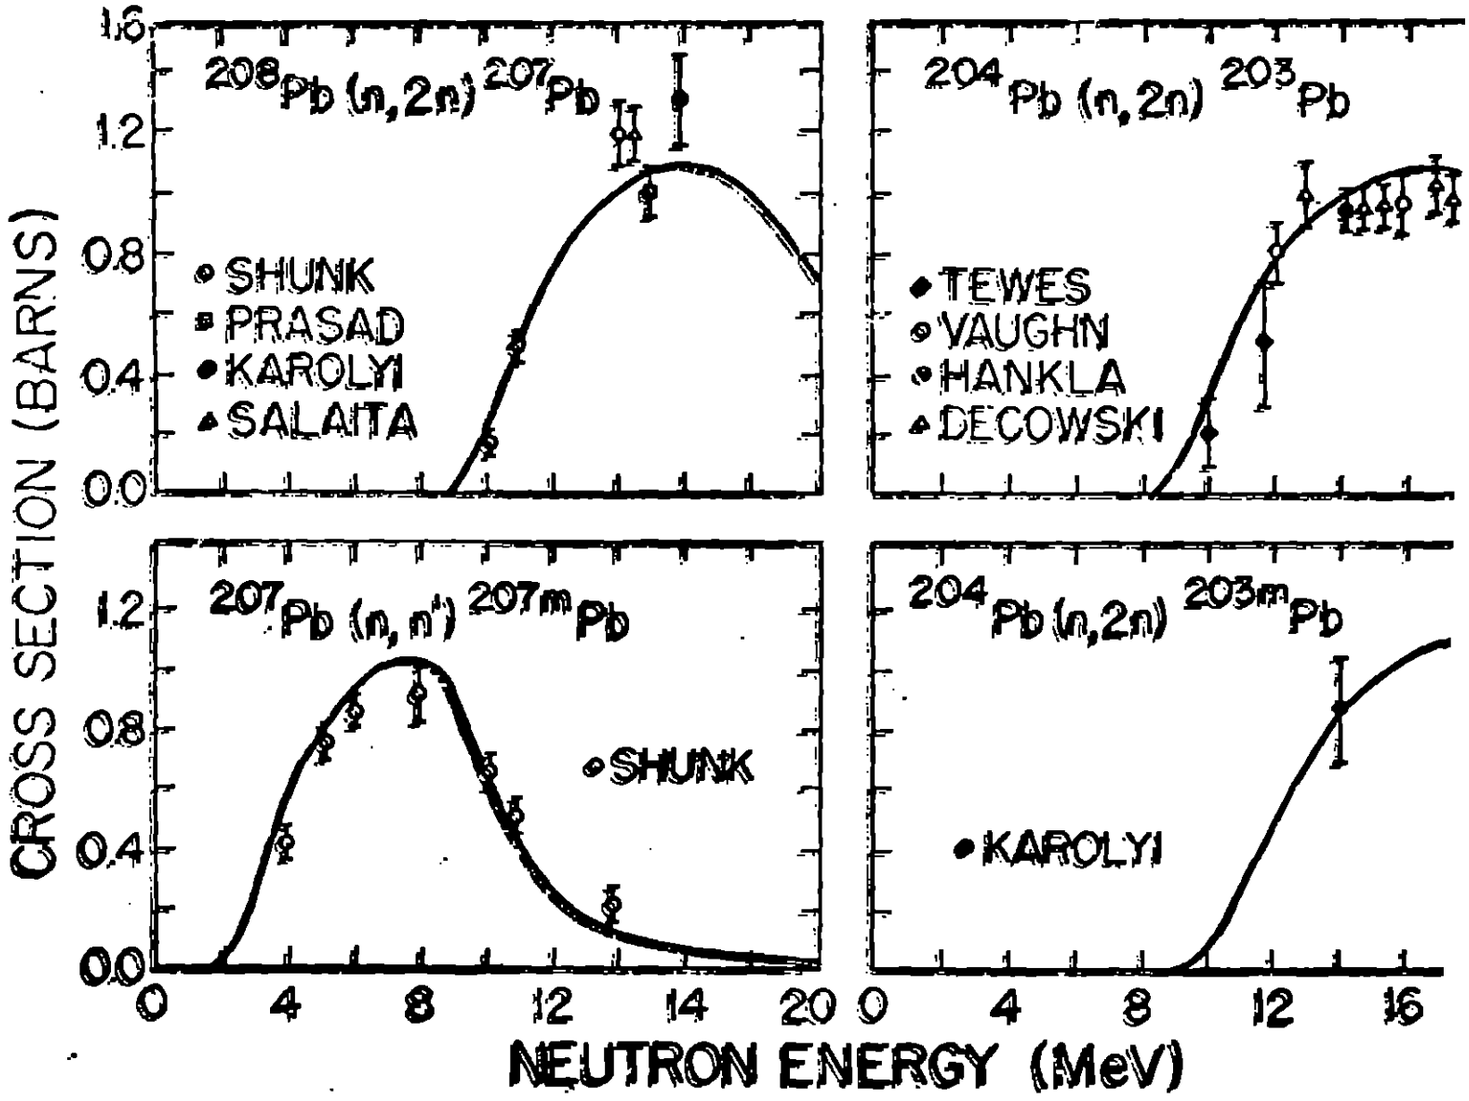
^{208}Pb (n, 2n) ^{207}Pb
SHUNK
PRASAD
KAROLYI
SALAITA
^{204}Pb (n, 2n) ^{203}Pb
TEWES
VAUGHN
HANKLA
DECOWSKI
^{207}Pb (n, n') ^{207m}Pb
SHUNK
^{204}Pb (n, 2n) ^{203m}Pb
KAROLYI
CROSS SECTION (BARNS)
NEUTRON ENERGY (MeV)
0.0
0.4
0.8
1.2
1.6
0
4
8
12
14
16
20