

# RELIABILITY EXPERIENCE WITH CHI/OS

W. C. Lynch

**Reporters:** Dr. S. Tsur  
Dr. I. Mitrani  
Dr. J. H. Hine

## Abstract

Professor Lynch continued the theme of his first lecture by discussing the reliability aspects of system performance. In particular, he drew attention to various aspects of the design process which influenced reliability, some a planned part of the design process and others arising from the environment in which the design was carried out.

## Introduction

CHI/OS is a system designed for the Univac 1108 at the Chi Corporation. Chi Corporation is a service bureau owned by Case Western Reserve University and supplying computing services for about 3500 customers both within and primarily without the university. The system was developed over a five-year period ending with its installation in November 1973. The development proceeded at a low level of effort over the entire period. Reliability was a prime consideration during the design and implementation of CHI/OS. When the system became operational, all files had to be converted to a new format, program libraries altered, etc. Thus, it would have been impractical to frequently retreat to the old system if CHI/OS had been found to be unreliable. Observations made during the design and development of CHI/OS suggest that environmental factors also contributed to the reliability and performance of the final system. This lecture discussed these factors. The contribution of any single factor must be a subjective assessment aided by hindsight. It is nearly impossible to repeat a controlled experiment, and further, the expense of each data point is prohibitive.

Factors which contribute to reliable software may be viewed under two headings, system structure and environmental factors. The system structure will be determined by choosing an appropriate design methodology such as structured programming, structured analysis, Petri nets, etc. CHI/OS is organized around structures comparable to Hoare monitors. The system makes extensive use of the P and V operators. Although the use of a good design methodology leading to an appropriate system structure was a necessary condition for reliability, it was not sufficient.

Environmental factors arise from the organization of the design and development of the system and the environment in which it is carried out. Two major factors were important in the organization of the CHI/OS project. First, there was a large degree of interaction among the various groups concerned. The design and development groups shared personnel in common. This led to the implementers having a better than average knowledge of the overall design. This contributed to the overall success of CHI/OS. There was also interaction between the development and operations staffs which contributed to the reliability. The second major environmental factor was the use of a systems implementation language, CHILI. The vast bulk of CHI/OS was written in CHILI, although the kernel was written in assembly language. Efficiency was a serious consideration in writing the kernel. The system executes 8000 P and V operations per second. The process structure imposed by the CPU was also written into the kernel rather than the system implementation language. These factors are considered in more detail after the workload and reliability of CHI/OS are surveyed.

## Workload

CHI/OS supports a commercial computer utility with a high percentage of batch and remote batch jobs and a relatively low volume of interactive work. It is run as a vigorous open shop with users submitting their own decks and removing their output. Operators intervene only when necessary, e.g. tape mount requests from remote sites, system crashes, etc.

A detailed breakdown of the workload on March 1975 is shown in Table I. It can be seen

that a large amount of system activity is devoted to spooling. Also, despite a good file system, there is a large amount of tape activity. Under these loading conditions, about 1% of available time is lost to software errors.

**35000 jobs/22-day month**

**75% commercial and CHI contract programming**

**25% CWRU academic computing and administrative data processing**

**I/O activity**

**Spooling**

- 22 high-speed RJE terminals**
- 10 low-speed (up to 1200 baud)**
- 30 lines printed/second**
- 10 cards read/second**

**Tape**

- 6\*10<sup>5</sup> bytes/minute**
- 1 mount/minute**

**Three-level virtual memory file system**

- 6 disk accesses/second**
- 50 drum accesses/second**

**Table I: CHI/OS Workload on March 1975**

**Reliability**

Table 2 shows the amount of time lost due to software failures and due to hardware failures. Reasonable success was attained in distinguishing these failures. It should be noted that after the first full month of operation, the percentage of time lost due to software errors was less than that due to hardware faults. Some further improvement was obtained until software reliability was not a problem relative to the reliability of the hardware. This point was reached after nine months of operation.

There appears to be some correlation between the number of crashes due to software and hardware, especially in the early months of operation. There are several possible explanations for this. Contingency conditions for recovering from hardware faults may not have been thoroughly checked out. Hardware faults are usually not classified.

Hardware errors may appear as software errors, at least until the fault is analysed. This phenomenon is aided by the limited scope of hardware diagnostics. The only useful diagnostic test is the operation of the operating system.

Although the table suggests an equilibrium has been reached, occasional increases in the number of software crashes have occurred since March 1975. These have typically been due to the installation of new hardware or software and have only effected one month's statistics. This phenomenon will also contribute to the apparent correlation between hardware and software failures.

The remainder of this lecture was devoted to a discussion of the factors which contributed to the reliability of CHI/OS.

### **Project Organization**

In retrospect, the organization of the project contributed significantly to its reliability, although this organization was partly dictated by the economics of the Chi Corporation rather than carefully planned. The development of CHI/OS carried a relatively low priority and, as a result, involved a small number of people over a long period.

One intentional decision was to separate some system designers from the day-to-day development. These designers were responsible for auditing the overall development and structure of the system. The intention was to avoid the problems which arise when each person on a project sees the design in his or her own context. The remaining designers and a development group faced the problems of day-to-day implementation. They were responsible for problems that could be solved by themselves after any appropriate discussions.

Over the course of the implementation of the system, eight people were involved in a part-time basis. At any given time, the total manning level was about three-fourths of a person. It was an important consideration to keep everyone involved in all aspects of the design of CHI/OS. Design decisions are often questions of philosophy and may be difficult to set down as explicit directives. By including the development group in discussions concerned with design, the flavor of the intention could be conveyed. The developers could then carry this through to the implementation. Many problems of interpretations of documentation can be avoided by ample discussion.

The project benefitted by having a small number of highly skilled people rather than a larger number of junior people. Although six to eight people were involved over five years, the total effort was about 7.5 manyears. This is lower than the norm which appears to be about 15 manyears in a university environment. It was suggested that people had more than the usual amount of elapsed time for their thought processes to operate, and this time may not have been accounted for. The size of the resultant system (30 cm. of listing) and the lower than normal effort would indicate that CHI/OS benefitted from this enlarged elapsed time. It should be noted that perhaps an additional 3 manyears was required to develop CHILI.

At the beginning of the development of CHI/OS, Chi Corporation underwent a quantum jump in computer power. The temporary surplus of machine time meant that the project had all that it required. Many vendors did not seem to encourage this in the development of their own systems in 1968. The computer time charged was twice the manhours charged to the project. Many development projects are not blessed with this good fortune. On the other hand, the time used would have been idle time had the project not gone forward.

System testing was carried out in two stages. As soon as a skeleton system was developed, it was used to carry out further system development. The use of top down design helped in developing the skeleton system. For example, the file system was designed and implemented before any disk driver was written, files being limited to drum storage. The second stage involved making the system available for users. The users were encouraged to try out various functions and provide feedback to the development group. CHI/OS had undergone two years of testing when it went operational in November 1973.

In response to a question about staff turnover, it was noted that only one person who had written a significant part of the system had left before it went operational. It was felt that a low turnover was important to the success of any project. This was aided by the close cooperation of the design and implementation people which produced a strong individual identification with a part of the system. People are not happy with too small a part in a

project.

## **Production**

CHI/OS was designed to facilitate rapid correction when a software error was detected. By keeping the development staff close to the operational environment, the user provided constructive feedback. By anticipating the occurrence of errors and planning for their correction, a system which stabilized quickly after large changes was achieved. Errors were identified and corrected rapidly, as opposed to significantly fewer errors occurring.

The typical modification sequence following the identification of an error is as follows. The difficulty is analyzed using source level tools which provide information on about 95% of the crashes. This analysis is done immediately and is often complete within one hour. The correction is made in the offending source module and the module recompiled. A systems generation requires only three to five minutes of elapsed time during the working day. This allows a new load tape to be generated as soon as the correction is made. The new system is loaded and tested overnight by the operations staff. (This insures good communication between the development and operations staff.) Because of this procedure, it is not unusual for an error to have been removed from the system the day after it is detected.

The ease with which a change to the system can be made has been convenient for system software monitoring. This is very much an interactive process. If one set of collected statistics indicate the need for a different measurement to be made, the new probe is inserted overnight and the statistic recorded on the following day. This same process was obviously very useful during the development process.

There was some question about the security of a system which students were able to modify so easily as part of their research. These students are typically closely involved with the system programmers and discuss their proposed changes or insertions. The same trust that is extended to the system programmers must also be extended to the students.

## **CHILI**

The system implementation language, CHILI, provided several facilities which contributed to the reliability of CHI/OS. In designing the language, as many decisions as possible were deferred. Descriptive methods were provided to handle register allocation and memory management. Excessive flexibility is not good. The ability to describe complex data structures and the fact that it was a high level language reduced the amount of code required. Perhaps the feature that contributed the most to reliability was the INCLUDE statement. This ensured that structure definitions and register allocations were consistent and that each procedure call had exactly the same format. These declarations could be put into the library once. CHILI still had several problem areas. Because memory management was deferred to the system, all temporary allocation in CHILI had to be specified explicitly. Also, there was a tendency on the part of the staff to feel that a high-level language was self-documenting. This was not true. It was also concluded that CHILI should have contained a facility for describing the process structure.

## **Conclusion**

Explicit considerations such as the use of CHILI, the inclusion of debugging tools and rapid systems generation contributed to the reliability of CHI/OS. Other factors which arose from the organization of the project also contributed. These included the relatively small staff and high degree of interaction of all people concerned.

Table 2: CHI/OS Reliability

## 1108 DOWN TIME DUE TO SOFTWARE FAILURES

| MONTH  | # OF<br>CRASHES | MINUTES<br>LOST | PERCENT OF<br>TIME LOST |
|--------|-----------------|-----------------|-------------------------|
| DEC 73 | 20              | 529             | 5                       |
| JAN 74 | 14              | 157             | 1                       |
| FEB 74 | 13              | 193             | 1                       |
| MAR 74 | 12              | 265             | 1                       |
| APR 74 | 9               | 79              | 1                       |
| MAY 74 | 9               | 61              | 0                       |
| JUN 74 | 9               | 143             | 1                       |
| JUL 74 | 8               | 336             | 2                       |
| AUG 74 | 7               | 81              | 1                       |
| SEP 74 | 6               | 47              | 0                       |
| OCT 74 | 6               | 53              | 0                       |
| NOV 74 | 2               | 9               | 0                       |
| DEC 74 | 6               | 30              | 0                       |
| JAN 75 | 6               | 56              | 0                       |
| FEB 75 | 7               | 53              | 0                       |
| MAR 75 | 7               | 57              | 0                       |
| TOTAL  | 134             | 2092            | 1                       |

## 1108 DOWN TIME NOT DUE TO SOFTWARE FAILURES

| MONTH  | # OF<br>CRASHES | MINUTES<br>LOST | PERCENT OF<br>TIME LOST |
|--------|-----------------|-----------------|-------------------------|
| DEC 73 | 14              | 171             | 2                       |
| JAN 74 | 15              | 179             | 1                       |
| FEB 74 | 11              | 166             | 1                       |
| MAR 74 | 6               | 185             | 1                       |
| APR 74 | 9               | 201             | 1                       |
| MAY 74 | 7               | 121             | 1                       |
| JUN 74 | 6               | 44              | 0                       |
| JUL 74 | 11              | 948             | 7                       |
| AUG 74 | 5               | 84              | 1                       |
| SEP 74 | 7               | 121             | 1                       |
| OCT 74 | 2               | 75              | 1                       |
| NOV 74 | 2               | 8               | 0                       |
| DEC 74 | 6               | 27              | 0                       |
| JAN 75 | 7               | 1587            | 11                      |
| FEB 75 | 7               | 248             | 2                       |
| MAR 75 | 14              | 481             | 3                       |
| TOTAL  | 115             | 4165            | 2                       |