AN EXPERT SYSTEM FOR PERFORMANCE EVALUATION OF UNIX-BASED COMPUTER SYSTEMS

Yong Meng TEO
Department of Information Systems
and Computer Science
National University of Singapore
Kent Ridge Road
SINGAPORE 0511
tcoym@iscs.nus.sg

Abstract

The design of computer systems has historically been focused on determining the needs of users and packaging the needs into applications. It is generally assumed that the computer system configuration and system software will deliver acceptable performance levels. Issues related to system performance and capacity have received insufficient attention even though its importance has been recognised.

Expert system technology offers potential to support performance evaluation and capacity planning of computer systems. This paper describes our development efforts in the design, development, and use of a cost-effective PC-based expert system, CPEES, for performance evaluation of Unix-based computer systems. We discuss in this paper the external conceptual architecture and internal structure of the proposed expert system, its prototype implementation on a Unix multiprocessor computer system, and the operation of CPEES and steps involved in the flow of consultation.

Keywords: performance monitoring, performance evaluation, expert system, capacity management

1 Introduction

With the increasing complexity and proliferation of computer systems, the problem of performance evaluation becomes more difficult [4, 5]. A number of tools exist to collect and report performance data. This relies on the user to examine the data collected and compare them with expected values to identify a performance problem and propose solutions. However, it requires a good deal of expertise and experience to interpret the potentially overwhelming amount of information that performance monitoring tools can provide. With a limited number of experts in a well defined area of expertise, the situation and problem domain is appropriate for an expert system.

Indeed, a number of expert systems for computer performance evaluation have been proposed or described. Some examples of efforts include YSCOPE, a shell for building expert system for an IBM/MVS environment [3]; ESP, a prototype expert system for IBM/MVS systems [6]; XCON, an expert system for configuring DEC computer systems [8]; a PROLOG-based expert system for tuning MVS/XA [2]. Each of these sytems has a slightly different objective, and requires a different amount of user expertise and involvement.

Our experience and illustrations come from the design and prototype implementation of a PC-based expert system for monitoring and performance evaluation of a multiprocessor system. However, the approach use should be appropriate for other performance evaluation environments as well. The paper is structure as follows. An overview of the performance problem together with the configuration of the machine used in prototyping is presented. The conceptual structure and internal organisation of our proposed system, Computer Performance Evaluation Expert System (CPEES), are discussed. The operation and flow of consultation are presented using examples. Lastly, the potential benefits of CPEES and some issues and problems are discussed.

2 Expert System for Performance Evaluation

A good knowledge of the tools available and the variables to measure is required for performance measurement, whereas good understanding of the object computer system and its workload is necessary for performance monitoring, system tuning, and performance evaluation. On the other hand, capacity planning and management require requisite skill to forecast future workloads on the system and trends in computer technology. Expert system offers a number of attractive features. The cost of providing expertise is greatly reduced especially when it is difficult and expensive to hire someone due to shortage of expertise in this field. The expert system's knowledge will last indefinitely. Expert system is the mass production of expertise. Through continuous enhancement, the knowledge base of the expert system may be able to handle 80 to 90 percent of the possible problems in a computer system. It can explicitly explain in some detail the reasoning that led to a conclusion. It may response faster and be more readily available than a human expert.

The particular performance evaluation problem that we are addressing is off-line bottleneck detection. Performance measures of a system collected using software monitor together with performance expertise (knowledge) are used as input in the analysis. The method adopted is consistent with that described in [4]. The tool, CPEES, simulate and automate the activity of the performance consultant through an interactive computer session. The problem is that of identifying the bottleneck component in the system, determine its cause and to recommend corrective action to remove it.

The system is developed on an IBM compatible microcomputer, using a rule-based expert system shell called Personal Consultant Plus (PCP Plus) from Texas Instruments [7]. Rules can be coded in LISP or in a language called Abbreviated Rule Language (ARL). PC Plus provides forward and backward chaining mechanisms, and supports frame-based knowledge representation.

2.1 Performance Problem

A multiprocessor AT&T Unix based system is used in the study. The system consists of four processors interconnected via a high-speed packet bus as shown in figure 1. A Master Processor (MP) runs the UNIX operating system, controls the other processors, assigns processes to the other processors, and manages the system's global resources. Two Enhanced Adjuncts Data Processors (EADPs) act as the computation and file servers. An Adjunct Communication Processor (ACP) besides performing computation also acts as a communication processor. By default, the operating system accomplishes automatic load balancing across these processors by assigning new processes to the least busy processor. To maintain proper system

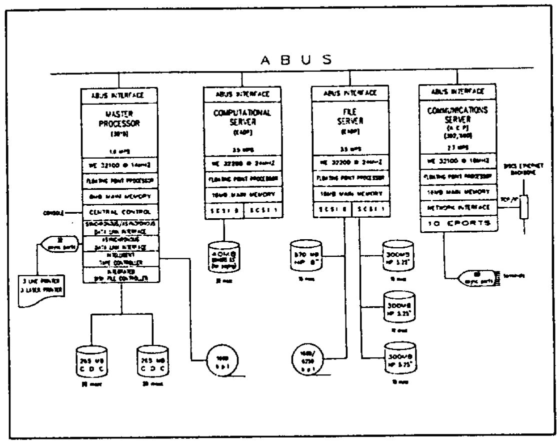

Figure 1: Multiprocessor System Configuration

operations, there are several aspects that have to be considered [1]. These include balancing the processor load, managing the file systems, balancing disk I/O, managing the bus traffic, and ensuring good response time. A poorly planned and managed system can lead to loss efficiency and effectiveness of the whole organisation.

2.2 External Structure

The external structure of the proposed system, CPEES, is shown in figure 2. A software monitor System Activity Report (SAR) collects performance data by writing pre-selected data into a log file at regular intervals during the collect period. A variety of reports can be generated covering times within the range of the collection period. Data extraction programs written in Turbo Pascal are use to extract relevant performance data for storing into a *performance database*. Knowledge in an expert system consists of a combination of facts and heuristics. The facts constitute a body of information that is widely shared, publicly available, and agreed upon by experts in the field. The heuristics are mostly private, little discussed rules of good judgement that characterise expert level decision making in the field. The performance level of an expert system primarily depends on the size and the quality of the knowledge base. In our prototype, expert knowledge and heuristics of the system for setting up the *knowledge base* are obtain mainly from human expert (i.e. systems programmer) and system manuals provided by the vendor. The prototype knowledge base created contains facts and rules for the four processors' CPU, memory and I/O resources. Based on monitored performance measures from the performance database, an *inference engine* executes relevant

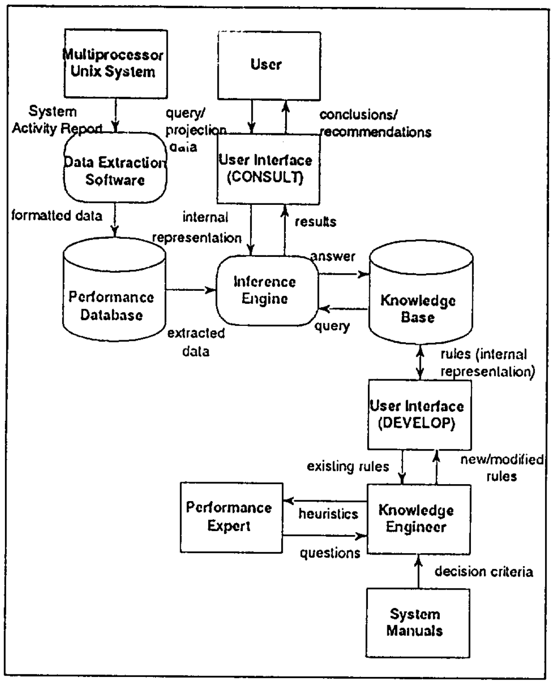

Figure 2: External Structure of CPEES

rules in the knowledge base and present to the user the conclusions/recommendations with explanations.

Two modes of operations are provided by CPEES, namely the consult mode and the develop mode. The *consult mode* provides consultation facility, and communicates with users the findings, recommendation and explanations. The *develop mode* facilitates the creation and maintainence of the knowledge base. The knowledge engineer uses this mode to translate the heuristics gathered from performance experts and manuals for storage in the knowledge base. This include coding of rules and modularisation of knowledge into frames. New rules can be added to the knowledge base as and when new information is gathered.

2.3 Internal Structure

A computer system consists of three main components: CPU, memory and I/O. Whenever any of these components limits system performance a performance bottleneck results. The internal structure of CPEES follows this classification. For each processor, any one of the three components can be selected for evaluation. This also enables us to organise the rules in the knowledge base into subframes as shown in figure 3.

The three subrames are:

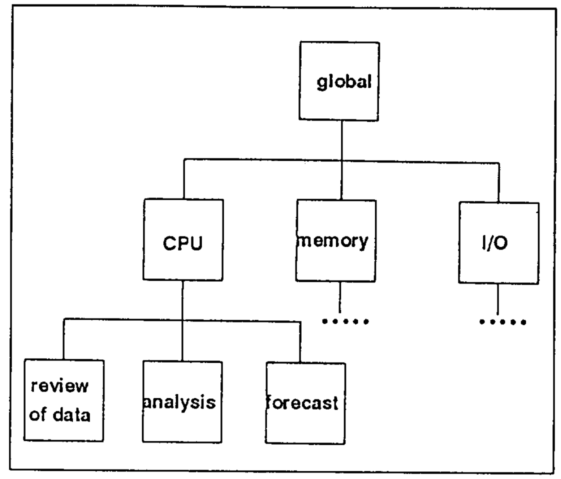

Figure 3: Internal Structure of CPEES

1. *CPU Resource Subframe:* The CPU in the system performs both user and system tasks. Since the CPU is a limited resource, these tasks must be balanced so that a sufficient throughput is maintained. The factors that can cause a decrease in CPU throughput include: users attempt to execute more processes than it can handle; memory resource limitations resulting in the CPU spending excessive amount of time in performing memory management operations.

2. *Memory Resource Subframe:* Each processor on the system has its own memory, and a paging disk. The memory requirement on each processor depends on the type of processes running on the processor, and different processors may exhibit different degree of paging and/or swapping activity. Some factors that leads to memory resource problems include: too many processes or large processes, poor locality of reference of applications.

3. *I/O Resource Subframe:* I/O resource limitations can cause a system to perform poorly. Some of the causes include: too many accesses being generated for the number and type of devices; improper file system balancing, resulting in certain disks being overloaded while others are idle; improper file system layout on the disk, causing increased seeks on the disks; allocation of swap areas on heavily used disks.

Since there are four processors in the multiprocessor system under evaluation, this translate into twelve evaluation frames. For each subframe, three modes of consultation facility namely *review*, *analyse* and *forecast* are provided. At the top level, a global (root) frame performs calling functions to instantiate each subframe to carry out analysis on the performance date. The global frame also consolidate the results and analysis from different subframes and outputs the overall conclusion/recommendation to the user.

Apart from the hardware configuration of a computer system, one should note that a performance bottleneck may depends upon the workload as well. Workload is currently not included in CPEES, but can be easily incorporated. Performance data that characterise workload can be obtained from the system accounting log file. If the workload is not sufficiently heavy, a performance bottleneck may not appear even though the configuration is suboptimal. Morever, because of the transient nature of the workload, bottlenecks may come and go over time so that different and contradictory recommendations may result from analysing different time segments within performance data collected. Although bottleneck detection does not address the entire performance optimisation problem, it does address the most important class of problems, namely when computer performance is actually limited.

2.4 Operations and Flow of Consultation

There are two modes of operations in CPEES: develop and consult mode. The develop mode enables the user to create and maintain the knowledge base. CPEES provides several menu and help screens to guide the user to modify, add and delete rules and preset parameters in the knowledge based. Figure 4 shows examples of rules for analysis CPU, memory and I/O. At the time of writing, the knowledge based contains about 200 hundred rules. This consists of 50 global rules for analysing the overall system, and 60, 40, 50 rules for analysing inidividual CPU, memory and I/O subsystems respectively.

The consult mode presents to the user the analysis, findings, recommendations and explanations and forecasting on the capacity of the various subsystems. After the extraction of performance data is completed, a user can select the review option to preview the performance data in textual or in graphical form. Six options provided include a help menu, a list of performance measures with explanations, preview of overall system performance and individual CPU (data are further grouped under CPU, memory, and I/O), data

```
RULE010
SUBJECT:: EVAL-CPU-80-RULES
DESCRIPTION:: (Excessive system calls occurring.)
IF:: (((SCALL/S-I-80+SCALL/S-O-80)+SCALL/S-L-80)>300)
THEN::
(ADVICE-C.P.U.-80=(TEXTEVAL(TEXTNAME CPU_010))
AND ADVICE-C.P.U.-80=((SCALL/S-I-80+SCALL/S-O-80)+
SCALL/S-L-80))

RULE203
SUBJECT:: FORECAST-CPU-121-RULES
EXPLANATION:: ("This rule calculates the estimated CPU
life whenthe workload increases.")
IF:: (FC-C.P.U.-DATA-121!=0 AND %IDLE-121≥10)
THEN::
(FC-C.P.U.-LIFE-121=(TEXTVAL(TEXTNAME FC_CPU_002))
AND FC-C.P.U.-LIFE-121=(FIX((((%IDLE-121 - 10)/(100 -
%IDLE-121))*1200)/FC-C.P.U.-DATA-121)) AND
FC-C.P.U.-LIFE-121 = "MONTH/S")

RULE073
SUBJECT:: EVAL-MEMORY-120-RULES
IF:: (((LEXEC/S-120+EEXEC/S-120)+IEXEC/S-120) > 2)
THEN::
(ADVICE-MEMORY-120=(TEXTVAL(TEXTNAME MEM_001))
AND
ADVICE-MEMORY-120=(VALUE LEXEC/S-120) AND
ADVICE-MEMORY-120=(TEXTEVAL(TEXTNAME
MEM_001_1)) AND
ADVICE-MEMORY-120=(VALUE EEXEC/S-120) AND
ADVICE-MEMORY-120=(TEXTVAL(TEXTNAME MEM_001_2))
AND
ADVICE-MEMORY-120 = (VALUE IEXEC/S-120) AND
ADVICE-MEMORY-120=(TEXTVAL(TEXTNAME
MEM_001_3)))

RULE028
SUBJECT:: EVAL-IO-80-RULES
IF::(%RCACHE-L-80<90 OR %WCACHE-L-80<60)
THEN::
(ADVICE-I/O-80=(TEXTVAL(TEXTNAME IO_006)) AND
ADVICE-I/O-80= (VALUE %RCACHE-L-80) AND
ADVICE-I/O-80=(TEXTVAL(TEXTNAME IO_006_1)) AND
ADVICE-I/O-80=(VALUE%WCACHE-L-80) AND
ADVICE-I/O-80=(TEXTVAL(TEXTNAME IO_006_2)))
```

Figure 4: Examples of Rules

visualisation using bar charts, pie-charts. For textual data, colour are used to highlight hot/cold spots when performance measures exceed predefined threshold values. After previewing the data, the user can proceed to analyse the overall system or to analyse selected system component. In analysing the overall system, the inference engine triggers the rules for all twelve system components sequentially. Result of the analysis is presented in three parts: *analysis* with explanations, *figures* showing the parameters used, and *recommendations* if any. Samples of consultation result screens are shown in figure 5. When the problem

Analysis:
The processor may be CPU limited. More than 90% of the time, the run queue had at least one process ready to run but was waiting for the CPU.
Figures:
% of time the run queue waited for CPU (%runocc): 92%
Recommendations:
Reschedule jobs to non-prime time.

Analysis:
There are too many processes being invoked per second.
Figures:
number of local exec system calls per second (lexec/s): 3.81
number of emigrate exec system calls per second (eexec/s): 3.38
number of immigrate exec system calls per second (iexec/s): 3.33
Recommendations:
Modify applications to reduce the number of process invoked.

This is the status of the CPU of the File Processor, based on figures called by SAR:
Growth potential for the CPU is: 55%
Estimated life before the CPU power is exhausted: 7 months

Analysis:
The processor is probably spending too much time waiting for disk jobs complete. This conclusion is based on the fact that the read buffer cache hit ration (%rcache) falls below 90 or the write buffer cache hit ratio (%wcache) falls below 60.
Figures:
read buffer cache hit ratio (%rcache): 99
write buffer cache hit ratio (%wcache): 0
Recommendations:
Educate users who open files needlessly or create excessively. Increase size of cache: NBUF and NHBUF.

Figure 5: Examples of Consultation

identified is related to another part of the computer system, an analysis on that relevant system component is performed automatically.

3 Conclusions

We have discussed in this paper the application of expert system for off-line evaluation of computer performance. The tool allows the permanent capturing of expert knowledge on computer performance. CPEES can be used to help the less experienced performance analyst to perform their job more proficiently, and also for the training new analysts. With more rules acquired and added to the knowledge base over time, a more indepth performance analysis can be made.

Development in expert system technology seems to offer potential and opportunities for exploitation in evaluating computer performance. It is envisaged that in future, most medium and large size computer installation will have some forms of built-in evaluation and monitoring tools to constantly check for the proper functioning of the system components as well as anticipating capacity expansion needs in advance. Expert system is likely to be one of those tools. Expert system can help to focus the human experts on real issues, and upgrade the level of capability of the novice to perform at an expert level.

References

[1] *AT&T 3B4000 & 3B15 Computers, Unix System V, Operations and Administration Series - Performance Management*, AT&T, 1988.

[2] B. Domanski, *A PROLOG-based Expert System for Tunig MVS/XA*, Procs of CMG '87 International Conference on the Management and Performance Evaluation of Computer Systems, pp. 160-166, USA, 1987.

[3] J. Hellerstein and H.V. Woerkem, *YSCOPE: A Shell for Building Expert Systems for Solving Computer Performance Problems*, Procs of CMG '85 International Conference on the Management and Performance Evaluation of Computer Systems, pp. 170-180, USA, 1985.

[4] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation and Modelling*, John-Wiley, 1991.

[5] K. Kant, *Introduction to Computer System Performance Evaluation*, McGraw-Hill, 1992.

[6] A.P. Levine, *ESP: an Expert System for Computer Performance Management*, Procs of CMG '85 International Conference on the Management and Performance Evaluation of Computer Systems, pp. 9-13, USA, 1985.

[7] *Personal Consultant Plus - Reference Guide*, Texas Instruments Incorporated, 1987.

[8] *The Artificial Intelligence Experience: An Introduction*, edited by S.J. Scown, Digital Equipment Corporation, 1985.