Robert Young

ARTHUR - JOB and PROCESS Monitor for CX and Pre-Cx Machines

I must start, of course, by offering nearly all of you here congratulations on the 200th anniversary of the out-break of organized crime. It is organized crime as far as I am concerned because your ancestors were rebelling against my monarchy. From the tone of the meeting, it would appear that the spirit of the Revolution has not died but, instead of it being directed at George III, the imperial power, remote and majestical, is now Hewlett-Packard.

The origins of ARTHUR, like King Arthur himself, lie back in the mythology. The HP 3000 was a mythological sort of machine to us when we first got it. It had stacks and segments and a virtual memory and was a whole heap bigger than anything any of us had ever tackled at first hand. In 1973, when we took delivery of the machine, it came into our Institute as a mythological beast of unknown and awesome temperament. The staff of the Institute were mostly traditional chemists, that is they used the techniques of wet chemistry. In 1970 a new Director arrived who introduced more physical methods and it was he who was behind the selection of the HP 3000 for the Institute. Of the 600 scientific and technical staff at the Institute, there were no more than 10 who had any experience of computing of any kind at all at that time.

In those early days we used MPE 'B' and for the sort of work our users were doing, this was quite adequate. However, as time went on and we trained the staff, that is the Institute staff not the computing staff, these naïve users became increasingly confident and adventurous in their use of system resources. The first real bottle-neck was the absence of spooling. We overcame this by prevailing successfully on Hewlett Packard to supply us with a pre-release version of MPE 'C'. It was around then that our users seemed to gain a lot of confidence in a very short period and we quickly ran into trouble with other system limitations.

Repeated attempts to get Hewlett Packard to produce the necessary system performance monitoring software were of no avail and it was with considerable reluctance that we decided to do this work ourselves. Fortune smiled upon us and we were able to take on a new member of staff in the Computer Laboratory.

He came straight to us from University where he had been studying chemistry. He had been doing some work using ICL equipment but it was only as an adjunct to his chemistry; by no stretch of the imagination was John Sowden a systems programmer. The other staff of the Laboratory were, of course, committed very heavily to the existing work programme and Dr. MacIntyre, the Head of the Laboratory at that time, dropped the problem into the lap of this new lad. It is to his considerable credit that he was able to boot-strap himself into the position of producing ARTHUR in 9 months.

Hewlett-Packard people were only of use to us at the casual advice level. We were able to do the work only because we had already procured some documentation (by a process I shall not mention here): a book of system tables, a 'bedsheet' and a microfiche set, all of MPE 'B'. By using these three sets of documents and a considerable amount of native wit we were able to steer John through the early stages and then he was able to fly on his own and produce the necessary goods.

ARTHUR appeared in two stages. The first one monitored the system tables and produced a print-out on the consol of CPU times and stack sizes for any in-core processes that ARTHUR had picked up. The work we really wanted from ARTHUR was for him to forceably re-schedule work which was being run from a session and consequently in the CQ on our machine down to the DQ where we run JOBS, when that was appropriate. It is a characteristic of teaching a large, näive user base how to use the machine with the power of the 3000 that the users' programmes get bigger and bigger and bigger, so that without them being too aware of it there comes a point where sessional users are using stacks of 20k or more. Efficiency of computation is not something you can introudce to the näive user when he first closes with the machine. The introduction of the idea of computation efficiency has to be delayed until such time as the man himself becomes concerned about improving his throughput. Only then is he prepared to learn more about the details of the use fo the machine; prior to that such information would be considered by him to be useless.

We released information whereby the sort of stack size and CPU time that we considered to be legitimate for sessional work were explained, and by the aid of ARTHUR we were able to have the results monitored. With the first edition of ARTHUR we tapped people on the shoulder to have the ordinances enforced, but when the final version of ARTHUR became available, it became quite automatic.

Re-scheduling is done on the basis of a Process since it is Processes that take time and absorb resources. It is quite obvious that for a Process to have any existence at all there must be means within the operating system to set them up and destroy them. What we did not know was if it would be possible to re-schedule a Process once it had been set up by the operating system. A good look at the privileged intrinsics in the microfiche set resulted in us finding REMOVE and INSERT which do just what their names suggest.

A copy of ARTHUR on magnetic tape, and the listing of the software, written in SPL are lying on the table outside and anyone who may wish to inspect it or have a xerox copy may do so. The tape is intended for the User Group Library but anyone who would like a copy before this becomes available should send a tape to me at Mill Hill directly; we can only handle 800 BPI.

ARTHUR runs as a JOB and makes use of GETPRIVMODE and so does not appear flagged as PRIV in SYSTEM DUMPS - HP shrug off PRIV USER mode DUMPS. It only takes a few tens of seconds CPU time in 24 hours so is not heavy on resources. You talk to it via the switch register and the console. The programme is very well documented: I claim that as being my particular contribution, there is hardly a line without a comment. (At this point the meeting applauded).

In the IBM Systems Journal number 4, 1975, there is an article entitled "Performance Analysis of Virtual Memory Time Sharing Sytems": dear HP, can we please have a similar article about the 3000 and NOT JUST FOR THE SERIES TWO MACHINES.

There is a useful utility not supported by Hewlett Packard called OVERLORD. Like ARTHUR, it runs as a JOB. Please, HP, could not ARTHUR and OVERLORD be

taken up and become a composite, supported utility? This would improve the overall use of the machine without asking HP to make alterations in the operating system with unknown and indeterminate effects throughout the system.

Up the Revolution.