Correspondence

Dear Sir,

Human Performance in Interactive Graphics Operations
[**R. A. Reynolds**, *The Computer Journal* **26**, 93 (1983)]

I find Reynolds' note interesting but consider that there are two areas which are unfortunately not covered.

(a) It is not clear whether the results quoted show a statistically significant difference between sessions of up to one hour and those over one hour. A minor further point here is whether the data should have been grouped in whole hours or whether a finer analysis of length of session would have shown, for example, that very short sessions were atypical.

(b) As to the interpretation of the results the author mentions that some operator errors will be accepted as valid commands but suggests that the number of decisions taken in any session should be a good measure of operator performance. From personal experience in interactive non-graphics work, the number of mistakes increases after a certain length of session; this would tend to reduce the difference between less than one hour and longer sessions. In any case, it is presumably appropriate decisions that constitute a good measure of operator performance rather than total decisions; one would want to distinguish between the operators who made the same total number of decisions but significantly different numbers of good decisions.

Yours faithfully
G. PHILLPOTTS
Statistical Department
Home Office
Queen Anne's Gate
London SW1H 9AT
UK

R. A. Reynolds replies
I cannot agree that it is useful to distinguish between individual operator performances. The intention of the exercise was to establish the typical working pattern of an average operator in order to optimize rostering.

I must of course agree that mistaken decisions should not ideally be included in a measure of performance. However, in a modern interactive draughting system such as the one described there are normally several ways of carrying out the same operation, some of which are more efficient than others in certain circumstances. The process of detecting mistaken decisions is therefore not as straightforward as is implied. In an observation of a real-life situation on the scale of the one described, where over 140 000 commands were issued, it is probably impracticable.

A laboratory experiment in which a number of operators carried out similar tasks could of course be devised, but great care would have to be taken that these tasks did in fact reflect a realistic working pattern.

Further comments from G. Phillpotts
I have the following further observations, in the order in which Mr. Reynolds has replied.

(a) I suppose my original letter may not have been sufficiently clear as regards distinguishing between operators who had different productivity in terms of good decisions (last sentence of my letter). My point was that just as one would in day-to-day management wish to distinguish between operators on this basis, so should statistical measures of the aggregate performance of a number of operators try to take this into account.

(b) I agree this is a problem.

(c) It occurs to me that an experiment along the following lines might meet the combined requirements of realism and good measurement:

Divide the operators into experimental groups. Within each group, they would work the same working pattern with a fixed length of terminal session. The different groups would be assigned different session lengths. Incoming work would be randomly allocated to groups and the outcome would be the time taken to 'deal' with a piece of work.

Dear Sir,

The Tower of Hanoi as a Trivial Problem

I must take issue with M. C. Er when he calls the Tower of Hanoi problem 'intractable'.[1] When the problem is used in introductory texts on computer science, any difficulty which occurs is more likely to be because the trivial nature of the problem leads to a lack of appreciation of the power of recursion, rather than because of a lack of insight. At least recursion emphasizes the structure of the problem.

More importantly, although Er provides a non-recursive analysis of the problem, it is insufficient. He analyses the *moves* of the discs and not the *status* of the discs after each move.

If we tabulate for each disc which peg it resides on after each move the solution is obvious. Using zero based disc and peg numbering the peg on which disc j resides after move x is given by

$$P_j \equiv (-1)^j \left\lfloor \frac{x + 2^j}{2^{j+1}} \right\rfloor \pmod 3$$

where $\lfloor\ \rfloor$ is the floor function. All the properties of the Tower of Hanoi problem follow trivially.

Yours faithfully
R. J. HEARD
Programmer
Queensland Institute of Technology
G.P.O. Box 2434
Brisbane
Queensland
Australia, 4001

Reference

1. M. C. Er, A representation approach to the Tower of Hanoi problem. *The Computer Journal* **25**, 442 (1982).

Dear Sir,

System Prototyping

The idea of system prototyping proposed by Dearnley and Mayhew (*The Computer Journal*, **26**, 36 (1983)) seems to me very powerful. As the authors say, prototyping is widely used for industry systems, but not much for computing systems, yet a computer with modern software tools is ideal for modelling real-world situations, including other computing systems, and is re-usable. Further, the task of system modelling is much more congenial than writing paper specifications, which leads to higher productivity.

I would like to add support to system prototyping by describing two cases in which I have been involved, both of which worked very well.

In the first, some years ago, a model of a process involving parcels of random dimensions being packed into rectangular containers, clearly demonstrated the superiority of certain kinds of strategic packing over random placement, with very little software design effort and small computer running costs.

In the second case, more recently, a model of a system involving the reservation by computer users of terminals, communications channels, and interactive processes on several service computers was created using Fortran and IDMS on an ICL 2970 mainframe. The intersections of the five dimensions of user, time, terminal, channel and interactive process were represented by a permanently established network of interconnections in the IDMS database. A reservation was represented by marking appropriate nodes. Airline seat reservation is a simple subset of this problem.

In a system of this complexity, it is very difficult to anticipate the resulting performance, as the time to search along the various axes of the stored data depends critically upon details of the data structure. The reservation system required very short response times, or the time needed to do reservations would approach the total terminal time available.

The model also included a reservation dialogue by which the user would obtain information about unreserved facilities and