to DCS file

SCIENCE AND ENGINEERING RESEARCH COUNCIL
RUTHERFORD APPLETON LABORATORY

COMPUTING DIVISION

D I S T R I B U T E D C O M P U T I N G N O T E 5 2 3

VISITS
Notes on a vist to Dr I Mitrani University of Newcastle 20.11.81

issued by
D A Duce
24.11.81

DISTRIBUTION: R W Witty
F Chambers (Logica Ltd)
D A Duce
Miss G P Jones
Investigators/Mitrani

INTRODUCTION

The previous project looked at several themes on the performance of distributed systems:

1. under different scheduling algorithms.

2. reliability: what happens to multiprocessor systems when processors breakdown. What are the tradeoffs between single and multi processor systems? It turns out that there is an optimum number of processors to use for a given situation which may be arbitrarily large, but is usually quite small.

This work resulted in 4 or 5 papers.

Last year Isi and Peter King (RA) started looking at the performance of local area networks. They have now developed a model of an ethernet-like (CSMA) network. Isi worked on this with Erol Gelenbe over the summer. The model is approximate but seems to be fairly good - as validated by simulations. The model works best for a large (say 20) numbers of stations. This work has now been written up (see file). The model is very different to the initial model presented at the Cambridge Simulation and Modelling SIG March 1981.

Isi and Peter are now looking at modelling the Cambridge Ring. At the minipacket level the ring behaves like a shared server with feedback.

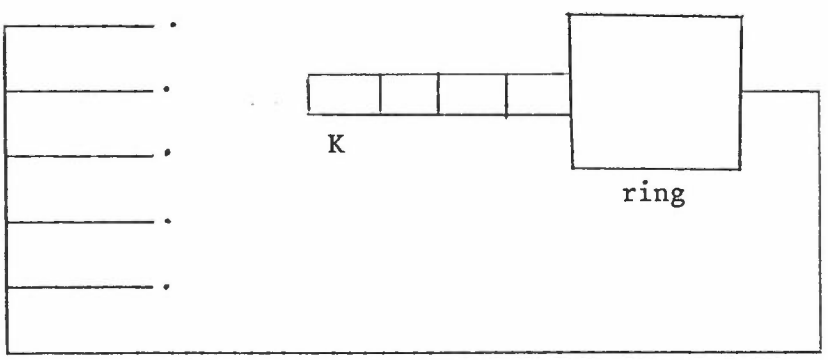

Stations submit requests to the ring. Assume there are k messages to be transmitted then shared-server proceeds at rate 1/(k+1), since empty minipacket cannot be re-used immediately

Message types are classified according to destination. Length distributions etc can be specified.

There is never any queuing at the receiving station. Basic block protocol using source select register prevents a station receiving other packets until the current block has been received completely.

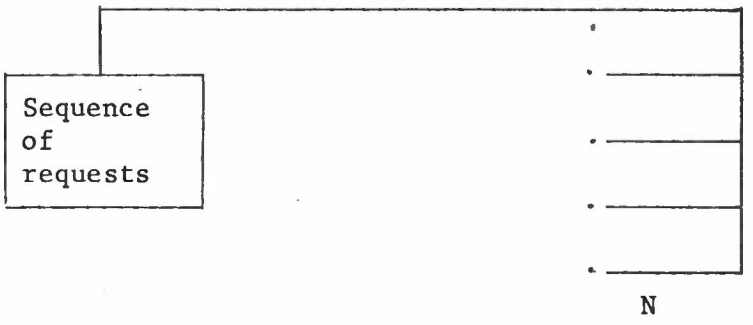

Nodes 1 to N are receivers. A request is of type i if it is directed to the i th receiver. There is a FIFO queue in front of each node. If K of the queues are not empty, K messages are being transmitted. The rate at which they proceed is 1/(K+1).

Classical theory does not apply to this model because the service rate is a function of the number of entries in the queue.

There are several possible approximations. One is being evaluated at present and compared with simulation results. The study is expected to take several months to complete. They hope to present the results at the SIGMETRICS conference in August 1982.

This model should be a basis for a comparison of ring and ethernet systems.

The effects of varying minipacket length can be investigated. In the limit this would represent a token ring - which some claim is superior to a slotted ring. However in a token ring a particular source destination pair can grab the ring for an arbitrarily long time interval.

Different access logics (DMA/PI etc) can be modelled by varying the service time per request.

When the local area network study is complete, they will move on to look at tasks executed on several processors.

Some structure for a job will be assumed, eg that tasks form a tree in which sons must be executed before fathers. Execution time will be studied as a function of the number of processors. If the tree is very flat, then many processors will not help, if it is bushy then there will be a noticeable improvement

Input to the model will consist of the probabilities that at any level in the tree a task has n descendants.

DISCUSSION

The coordinators stressed the value of prediction LAN models, especially when coupled with WAN models. Kent for example face the problem of how to network

a geographically very distributed site. Predictive models would be a real aid. Peter and Isi think their ring model will cope with bridges etc.

Unreliability has not yet been included in the models. The study of breakdowns in multiprocessors was none trivial. With 10 or 11 processors rounding errors influenced the results of a numerical implementation of the model. Breakdowns were modelled as a job with preemptive priority - this is one of the classical problems in queuing theory - it can be solved for some systems but not all.

The LAN models are only interesting with heavy workloads - much heavier than encountered in practice so far.

MOD (ASWE) have approached Isi - they want to look at performance of a network they are building. This is a very recent contact, made through Tom Anderson's project.

Isi would be happy to collaborate with anyone planning a network and wanting to model first. One problem seems to be that choice of configuration is made on some more or less "irrational" basis (eg cost!), rarely does an evaluation get to the stage where the choice is between A and B in performance.

The coordinator felt there is an isomorphism (some where!) between the multitask systems they propose to look at next and reduction/dataflow models and encouraged them to explore this avenue.