Analysis of locking policies in database management systems

D. Potier, P. Leblanc

► To cite this version:

D. Potier, P. Leblanc. Analysis of locking policies in database management systems. [Research Report] RR-0023, INRIA. 1980. inria-00076538

HAL Id: inria-00076538

<https://hal.inria.fr/inria-00076538>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

IRIA

Rapports de Recherche

N° 23

**ANALYSIS OF
LOCKING POLICIES
IN DATABASE
MANAGEMENT SYSTEMS**

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
B.P. 105 78150 Le Chesnay
France
Tél. 954 90 20

Dominique POTIER
Philippe LEBLANC

Juillet 1980

ANALYSIS OF LOCKING POLICIES
IN DATABASE MANAGEMENT SYSTEMS

Dominique POTIER
Philippe LEBLANC

* INRIA - Domaine de Voluceau - BP 105 - 78150 Le Chesnay - FRANCE

RESUME

L'accès concurrent de plusieurs transactions à un système de base de données impose la mise en place d'un mécanisme de contrôle de cohérence. Le contrôle est généralement réalisé en partageant la base de données en unités de verrouillage ou granules, et en définissant des règles de verrouillage des granules qui assure l'intégrité de l'information. Toutefois, le maintien de l'intégrité par ces mécanismes se fait au détriment des performances globales du système. Ceci est dû aux restrictions d'accès des transactions à la base imposée par les politiques de verrouillage, et à la charge supplémentaire induite par la gestion des verrous. Une approche d'analyse quantitative du rôle de ces facteurs est présentée dans cette étude. Dans une première étape, les facteurs les plus importants qui déterminent le comportement final du système du point de vue de ses performances sont mis en évidence et analysés indépendamment. Les résultats ainsi obtenus sont réunis dans une seconde étape pour conduire une évaluation des performances finale. Au cours de cette approche de modélisation hiérarchisée, diverses techniques de résolution analytique sont utilisées et les résultats obtenus sont illustrés par des exemples numériques. Les conclusions de l'article indiquent la sensibilité des résultats observés aux hypothèses concernant le comportement des transactions et la nécessité de données expérimentales dans ce domaine.

ANALYSIS OF LOCKING POLICIES IN DATABASE MANAGEMENT SYSTEMS

ABSTRACT

Consistency control has to be enforced in database management systems (DBMS) where several transactions may access concurrently the database. This control is usually achieved by dividing the database into locking units or granules, and by specifying a locking policy which ensures integrity of the information. However, a drawback of integrity enforcement through locking policies is the degradation of the global system performance. This is mainly due to the restriction imposed by the locking policies to the access of transactions to the database, and to the overheads involved with the management of locks. A framework for the quantitative analysis of the impact of these factors on the performance of DBMS is presented in this paper. In a first step, the main factors which determine the behavior of these systems are pointed out and analyzed independently. The results hereby obtained are aggregated in a second step to yield a global performance evaluation. Throughout this hierarchical modeling approach various analytical techniques are used and the result are illustrated by numerical examples. The paper concludes by pointing out final results sensitivity to some basic assumptions concerning transaction behavior and the need for more experimental studies in this area.

KEY WORDS AND PHRASES : consistency, concurrency, database management,
performance modeling, queueing models, queueing networks.

C.R. CATEGORIES : 3.50, 4.33, 8.1

I - INTRODUCTION

Quantitative analysis of locking mechanisms and of their impact on the performance of transaction systems have received relatively little attention [Ris77, Ris79, BIG79]. Although numerous concurrency mechanisms have been proposed and implemented, there is an obvious lack of experimental as well as analytical studies of their behavior and their influence on system performance. The present state of the art in this area can be compared to the situation in the late sixties when the first virtual memory systems were being implemented : at this time the now classical concepts of program behaviour (locality of references, working-sets, memory management policy, etc...) were only emerging and it took several years to fully understand and master them.

In order to perform a quantitative analysis of transaction systems of the performance, it is essential to point out the factors that determine the behavior of these systems. Three main factors may be identified :

i) transaction behavior : it may be described by the pattern of references of the transactions to the subsets (logical or physical) of the database and by the way the Central Processing Unit (CPU) and Input/Output (IO) resources of the system are used by the transactions. Given a locking mechanism, the reference pattern is obviously a determining factor : depending on the way references are distributed over the database, the locking mechanisms will have a more or less drastic effect. However, as noted above, little experimental evidence on reference patterns of transactions is available. In particular, the existence of properties of sequentiality or locality in the access pattern is still a controversial issue [RoH75, TuR76, Smi78, Rod76]. Faced with this situation, any analysis relies on rather strong assumptions. For instance, in the analysis of locking granularity presented in [Ris77], Ries and Stonebraker characterize the behavior of a transaction by the number of entities referenced, and assume that the number of granules locked by a transaction is proportionnal to the number of entities referenced.

In our analysis, we shall use a somewhat more sophisticated probabilistic model of transaction behavior although it also makes a strong assumption on the distribution of references over the database.

ii) locking scheme : locking schemes may be physical or logical depending on whether the units of locking are physical or logical subsets of the databases. We restrict our attention in this paper to physical locking. A physical locking mechanism is characterized by three main factors :

- the Lock Acquisition Policy (LAP) which specifies at which instant of the transaction lifetime locks are requested. LAP's are basically split into two categories: static LAP's (SLAP) and dynamic LAP's (DLAP). Under SLAP all the locks needed for executing transaction are requested at the initiation of the transaction. The transaction is executed only if all its requested locks are granted ; otherwise it is blocked. Under a DLAP, locks are requested by a transaction on demand, whenever a transaction references a granule that is not already locked.

- the Lock Release Policy (LRP) which specifies when the locks acquired by a transaction are released. In order to maintain the consistency of the database a static LRP (SLRP) has to be used, i.e. a LRP where a transaction releases its locks only when its execution has been completed.

Following [RiS77], the analysis presented in this paper is carried out by assuming that the locking policies used are SLAP and SLRP. The main advantages of this combination of LAP and LRP is that deadlocks are avoided and that, from the point of view of the analyst, its study is easier. However, it should be kept in mind that locks mean lifetimes are longer under a SLAP than under a DLAP, and hence that a SLAP restrict the transaction's accesses to the database more severely than a DLAP.

- the size of the physical locking units or granules (assuming that all locking units have the same size). Given a LAP and a LRP, the remaining

design parameter is the granule size. Obviously, in theory, the smaller the granules, the smaller the constraints imposed by the locking mechanism. However, at this point, secondary factors such as locking overheads have to be taken into account, since their influence may become important when the number of granules is large.

iii) the multiprogramming environment : the multiprogramming environment describes the resources (CPU, IO) that the transaction system uses for the actual processing of transactions. These resources, their characteristics and organization have an important impact on the global performance. Of the three factors we have identified, the latter is also presently the best understood and mastered. Within the past ten years much experimental and analytical effort has been devoted to this area [DeK76, MeP78], and the appropriate methods and tools for the analysis of the multiprogramming environment do exist.

Once these factors have been identified, we must select which approach will be used to analyze them and quantify their influence. In [RiS77, MuK77] simulation models are used. Although, this technique has the advantage of imposing no restriction on the complexity of the analyzed mechanisms, it cannot always provide a clear insight and understanding of the system performance and of the role of key assumptions.

Our approach in this paper is based on hierarchical analytical modeling [ChS78]. The basic idea underlying this approach is that in a first step each critical factor should be analyzed independently, clearly stating assumptions and their consequences ; the results hereby obtained are aggregated in a second step to yield a global performance evaluation and final conclusions.

A broad presentation of this approach is given in section II. Section III is devoted to the model of transaction behavior and the analysis of the locking mechanism. The multiprogramming environment is studied in

in section IV. Final results on the global performance are presented and discussed in section V. Conclusions are given in section VI.

II - PRESENTATION OF THE APPROACH

The transaction system is analyzed using a hierarchical modeling approach. Three levels of modeling are considered as represented in figure 1.

At level 1, the different stages transactions go through during their lifetime are described. The model used at this level is represented in figure 2. Transactions are issued from a set of interactive terminals. Upon arrival of a new transaction in the system all the locks needed for the complete execution of the transaction are requested. Lock requests are processed by the CPU and I/O resources of the system. During the "locks request" phase the transaction is in the state REQUEST. Upon completion of this phase, locks requested are either granted, either denied. If locks are granted, the transaction enters the state ACTIVE and it is executed by the CPU and I/O resources of the system until completion. It then returns to the terminal that has issued it. If locks are denied, the transaction is sent into the BLOCKED queue and enters the state BLOCKED.

BLOCKED transactions are removed from the BLOCKED queue in order to enter a new "locks request" phase when an ACTIVE transaction leaves the system and releases its locks. The maximum number of BLOCKED transactions removed from the BLOCKED queue at each departure is a control parameter of the transactions scheduling policy.

The organization and operations of the CPU and I/O resources of the system are described and analyzed at level 2. These resources are used by the transactions when they are either requesting locks (state REQUEST), or being executed (state ACTIVE). The analysis of level 2 will be conducted to derive the processing rates of ACTIVE transactions and REQUEST transactions depending on the number of these transactions sharing the CPU and I/O resources.

The behavior of transactions during their locks request phase will be analyzed at modeling level 3. Given a simple model of the transaction references requests to the entities of the database and the number of granules in the database, we will derive the probability distribution of the number of blocked granules depending on the number of ACTIVE transactions. These results will be used to obtain the probability that locks are granted to a transaction at the end of its locks request phase depending on the number of ACTIVE transactions.

The three levels of modeling will be presented and analyzed in the following order : Level 3, Level 2 and Level 1.

III - ANALYSIS OF TRANSACTION BEHAVIOUR AND LOCKING MECHANISMS

III.1 - The Physical database

The physical database consists of cells that represent the elementary physical unit of storage (this will be for example a disk sector) and it is divided into m granules, where the granule is the locking unit.

III.2 - Transaction behavior

The behavior of a transaction consists of a sequence of accesses to the cells of the database. This sequence is defined by the following assumptions :

- 1 - a transaction makes n accesses during its execution,
- 2 - accesses are uniformly distributed over the cells of the database,
- 3 - accesses are independent.

Given these assumptions we shall derive the probability $P_k(\ell)$, $\ell = 1, \dots, m$, that ℓ granules are locked when k transactions are simultaneously active, and the probability q_k that a new transaction enters the set of active transactions when k transactions are already active. These results will then be illustrated by numerical examples.

III.3 - Derivation of $P_k(\ell)$ and q_k

Let $x(p)$, $p = 1, \dots, m$ be the probability that the number of distinct granules accessed by a transaction is p , given that the transaction makes n accesses. The derivation of $x(p)$ is similar to a standard combinatoric problem : given n like objects and m unlike buckets, what is the probability of leaving p buckets non-empty after having uniformly distributed the objects into the m buckets ? Applying standard results of combinatorics (see for instance [Rio58] pp. 103), we then have

$$(1) \quad x(p) = \binom{m}{p} \binom{n-1}{p-1} / \binom{m+n-1}{n}$$

and, obviously

$$(2) \quad P_1(\ell) = x(\ell) \quad , \quad \ell = 1, \dots, m$$

In order to derive $P_k(\ell)$, we write the following recurrence formula between $P_k(\ell)$ and $P_{k-1}(\ell)$,

$$(3) \quad P_k(\ell) = \frac{1}{\sum_{j=1}^{m-1} P_{k-1}(j)} \sum_{i=1}^{m-1} P_{k-1}(i) \cdot \frac{x_i(\ell-i)}{\sum_{j=1}^{m-i} x_i(j)}$$

where

$$(4) \quad X_i(p) = x(p) \binom{m-i}{p} / \binom{m}{p} = \binom{m-i}{p} \binom{n-1}{p-1} / \binom{m+n-1}{n}$$

In (4), $X_i(p)$ is defined as the probability that a transaction is granted p granules given that i granules are already locked. The terms in the denominator of (3) are normalisation factors : since the k -th transaction can become active only if there are less than m granules locked by the $(k-1)$ st transactions, the probabilities $p_{k-1}(i)$, $i = 1, \dots, m-1$ are normalized by the sum $\sum_{j=1}^{m-1} p_{k-1}(j)$; in the same fashion, given i , the maximum number of new granules locked by the k -th transaction is $m-i$, so that the probabilities $X_i(n)$ in (3) are weighted by $\sum_{j=1}^{m-i} X_i(j)$.

From $P_k(\ell)$, we can derive the mean number $\bar{\ell}_k$ of granules locked when k transactions are active

$$(5) \quad \bar{\ell}_k = \sum_{\ell=1}^m \ell \cdot P_k(\ell)$$

and we straightforwardly obtain the probability q_k that a new transaction is granted the locks requested when k transactions are already active :

$$(6) \quad q_k = \sum_{i=1}^{m-1} P_k(i) \sum_{p=1}^{m-i} X_i(p)$$

III.4 - Discussions and numerical examples

The results obtained in the previous section may help us to have a first insight of the effect of locking on the operation of the transaction system. The mean number $\bar{\ell}_k$ of locked granules and the probability q_k give clear indications on the maximum possible number of simultaneously active transactions, and on the chance that a new transaction has its locks granted.

These indications will have to be analyzed and interpreted with respect to the model of transaction behavior that has been used. Recall that the model assumes that, for a given transaction, its accesses to the

cells of the database are uniformly distributed over the whole set of cells : in other words we assume that transaction references exhibit no locality. We may thus expect with this model of transactions behavior the effect of locking to be severe.

Those observations are illustrated by the numerical results presented in Tables I and II. These results have been obtained for $n = 2^i$, $i = 1, 5$, $m = 2^j$, $j = 1, 8$ and k ranging from 1 to 8. Table I and Table II respectively give $\bar{\ell}_k$ and q_k . Figures 3a and 3b depicts the variation of $\bar{\ell}_k$ as a function of m with $n = 4$ and $n = 16$.

The importance of the no-locality assumption is revealed by the influence of the number n of accesses made by a transaction on the mean number of blocked granules $\bar{\ell}_k$. Transactions with a large number of accesses to the database very rapidly tend to lock most of the granules and to prevent activation of new transactions. On the other hand, this effect is partly alleviated by an increase of the number of granules. However, as we shall see in the following sections, the increase of the number of granules causes an increase of the CPU and I/O activity devoted to locks requests, which is detrimental to the processing rate of active transactions.

IV - ORGANIZATION AND OPERATIONS OF THE CPU AND I/O RESOURCES

At a given instant, the transactions in states REQUEST or ACTIVE share the CPU and I/O resources of the system. We shall assume that there is no priority scheme set between ACTIVE and REQUEST transactions, and that the CPU processes these transactions on a processing-sharing (PS) fashion and I/O units process them according to a FIFO policy. The system consists of one Central Processing Unit (CPU) and a set of B identical independent Disk Units (DU). The I/O operations of a transaction are uniformly distributed over the B DU's.

We shall now describe the different CPU and I/O operations performed by REQUEST and ACTIVE transactions.

IV.1 - Operations performed by REQUEST transaction

Let $v(m)$ be the mean number of distinct granules requested by a transaction. A transaction in the state REQUEST will thus perform on the average $v(m)$ elementary lock request operations. An elementary lock request operation consists of :

- a CPU service time with mean t_R
- a DU service time with mean s_R .

An expression of $v(m)$ is simply obtained from the results presented in the previous section. We have

$$v(m) = \sum_{p=1}^m p \cdot x(p)$$

and, from (1),

$$(7) \quad v(m) = \frac{mn}{m+n-1}$$

It should be noted here that we have assumed that the elementary operations involved in a lock request, i.e. the CPU and DU services do not depend on m . This is certainly an optimistic assumption since management of locks tables may become more costly as m increases, and, for instance, several disk accesses may be needed per lock request for a large value of m . Therefore, results obtained for fixed values of t_R and s_R will have to be interpreted with respect to this assumption.

IV.2 - OPERATIONS PERFORMED BY ACTIVE TRANSACTIONS

A transaction in the ACTIVE state performs n accesses to the entities of the database, each access being followed by a CPU processing time. The activity of an ACTIVE transaction will thus consists of n elementary operations comprising :

- a CPU service time with mean t_A
- a DU service time with mean s_A .

We assume that the CPU and DU service times are independent random variables, and that DU service times have a negative exponential distribution.

IV.3 - Processing rates of ACTIVE and REQUEST transactions

Let k and p be respectively the number of ACTIVE transactions and REQUEST transactions processed by the resources at a given instant of time. Due to the fact that ACTIVE and REQUEST transactions have different behaviors their processing rates depend on k and p through functions of the form $\mu_A(k, p)$ and $\mu_R(k, p)$.

In order to derive $\mu_A(k, p)$ and $\mu_R(k, p)$, we follow the standard approach used for the throughput analysis of multiprogramming systems [DeK76]. As represented in figure 4, fixed numbers k and p of ACTIVE and REQUEST transactions are maintained in the system, and the sharing of the CPU and DU resources by the two types of transactions is analyzed using a closed queueing network model with different classes of customers [BaC75].

An analysis of this model will yield for instance the throughput rate $X_A(k, p)$ and $X_R(k, p)$ of ACTIVE and REQUEST transactions through station CPU. Since ACTIVE and REQUEST transactions require on the average respectively n and $v(m)$ CPU service times, we have

$$\mu_A(k, p) = X_A(k, p)/n ,$$

$$\mu_R(k, p) = X_R(k, p)/v(m).$$

The functions $\mu_A(k, p)$ and $\mu_R(k, p)$ aggregate the different effects of resource sharing by ACTIVE and REQUEST transactions. Using the equivalent server approach [ChS78], they will be used at modeling level 1 to analyse the global performance of the transaction system.

It should be noted that $\mu_A(k, p)$ and $\mu_R(k, p)$ depend on the granularity m through the function $v(m)$. Increasing the number of granules will increase the number of elementary lock requests performed by a transaction in the state REQUEST, and limit the availability of the CPU and DU resources to ACTIVE transaction. As a consequence, the processing rate of ACTIVE transactions will be degraded.

In order to illustrate this effect, we have plotted in figure 5 $\mu_R(k, p)$ as a function of p for m ranging from 2 to 2^8 with $n = 4$ and $k = 2$. The results have been derived with the following values of the parameters :

$$t_A = 50 \text{ msec.}$$

$$t_R = 10 \text{ msec.}$$

$$s_A = s_R = 50 \text{ msec.}$$

$$B = 3 \text{ (3 identical DU's)}$$

They have been obtained by using the solution package QNAP [MeP78, Ver79] developed by CII-Honeywell Bull, and INRIA. The listing of the QNAP program used for this computation and example of its outputs are given in Table III. Note that the value of $1/\mu_A(k, p)$ and $1/\mu_R(k, p)$ are stored in an array after each resolution so as to be used later on when global performance is analyzed.

It can be observed from figure 5 that, due to the specific behavior of the function $v(m)$, the effect of granularity on the locking overheads is particularly important as m increases from 2 to 2^5 .

V - ANALYSIS OF GLOBAL PERFORMANCE

Having analyzed modeling levels 2 and 3, we now dispose of the aggregate results that are needed for the analysis of the global performance. The interactions between the different factors which have already been investigated are summarized in Figure 6. Analysis will be conducted at modeling level 1 by using a multiclass queueing network model as presented in figure 7. The set of K terminals is modeled by an infinite servers station TERMINAL with mean thinking time T , and the set of CPU-I/O resources is replaced by two exponential FIFO stations RESREQ and RESACT, with processing dependent rates $\mu_A(k, p)$ ($\mu_R(k, p)$) when k ACTIVE transactions and p REQUEST transactions are respectively present in these stations.

When a REQUEST transaction has completed its lock requests, its locks are granted with probability q_k , if k transactions are active, and it then enters the set of ACTIVE transactions ; with probability $1-q_k$, its locks are denied and the transaction proceeds to the BLOCKED queue.

Transactions are removed from the BLOCKED queue when ACTIVE transactions terminate and release their locks. Thus removal of transactions from the BLOCKED queue is triggered by departures from the station RESACT. Let r be the maximum number of transactions removed from queue BLOCKED. If i is the current number of transactions in queue BLOCKED, the actual number of transactions removed from queue BLOCKED at each departure is $\min(i, r)$.

Due to the strong dependence between queues BLOCKED, RESACT and RESREQ, standard results on multiclass queueing network models are not longer applicable. Given the assumptions of the model, the only exact solution technique available is through a Markovian analysis [Ste78]. This is done by using again the solution package QNAP. The text of the QNAP program is given in Table IV. Transfers from the BLOCKED queue to the RESREQ queue are specified by the primitive MOVE. Since QNAP uses only linear arrays a macro-instruction IND(I) is used to access the element of the arrays TA and TR where $1/\mu_A(k,p)$ and $1/\mu_R(k,p)$ are stored.

The results are presented in figure 8 and 9. They have been obtained with the following values of the parameters of the model :

$$\begin{aligned} n &= 4, \\ m &= 2^0 \text{ to } 2^8, \\ t_A &= 50 \text{ msec.}, \\ t_R &= 10 \text{ msec.}, \\ s_A = s_R &= 50 \text{ msec.}, \\ B &= 3, \\ r &= 1, \\ T &= 1.5 \text{ sec.}, \end{aligned}$$

Figure 8 displays the mean number of ACTIVE, REQUEST and BLOCKED transactions in the system as a function of m . The main observation is that as m increases the mean number of BLOCKED transactions decreases rapidly, while the mean number of ACTIVE transactions steadily increases.

The mean number of REQUEST transactions reaches a maximum for $m=32$ and slowly decreases afterwards. It is thus to be expected that the global throughput will increase with m for large values of m .

This is illustrated in figure 9 where is plotted the throughput $D(m)$ of the system, for n ranging from 2 to 2^5 , expressed in terms of mean number of transactions processed per unit of time. It is observed that $D(m)$ decreases through a minimum and then increases with m .

This observation, which contradicts the results presented in [RiS77] where $D(m)$ exhibited a maximum, is to be interpreted with respect to the set of assumptions that have been used through the analysis. As it appears from the results of modeling levels 2 and 3, since q_k increases with m , the ratio of the number of ACTIVE transactions to the number of REQUEST transactions tends to increase ; on the other hand, since the mean number of locks $v(m)$ tested by REQUEST transaction increases with m , REQUEST transactions tend to need more processing as m increases. However, as indicated by the behavior of the throughput function $D(m)$, this behavior is determined by the increase of q_k rather than by the increase of $v(m)$. This points out again the key role played by locking overheads, here quantified by the function $v(m)$. In order to explain the discrepancies observed between our results and those presented in [RiS77], it should be recalled that the function $v(m)$ used in this study is of the form $v(m) = \alpha m$. In general, assumptions on transaction behavior used in those two studies are basically different : Ries and Stonebraker use a "linear" model of transaction behavior ; we use a "non-linear" model.

VI - CONCLUSIONS

We have presented in this paper an analytical framework for the performance analysis of locking mechanisms in transaction systems, and we have illustrated this approach by a detailed analysis based on a simple probabilistic model of transaction behavior.

This analysis provides a clear understanding of the various factors that determine global performance. It also raises many new issues that can only be solved by further extensive experimental and analytical studies. Two particular topics deserve special attention : the modeling of transaction behavior and the modeling of locking overheads. As noted above, the model of transaction behavior we have used makes a strong assumption on the distribution of references over the database : this model has to be validated by comparing its results to those of more complex models, involving for instance localities or sequentialities in the reference pattern. We have also pointed out in Section V the key role played by locking overheads. Here again a more refined analysis is needed, in relation with the modeling of transaction behavior. We are currently working along those directions and we hope that this paper will initiate parallel investigations.

ACKNOWLEDGEMENTS

The authors would like to thank J.L. COLOMER of the Facultad de Informatica, Universidad de Barcelona, SPAIN ; R. BALTER of the Centre Scientifique CII-Honeywell Bull in Grenoble, FRANCE ; and M. SCHOLL of INRIA for their helpfull contributions.

REFERENCES

- [BaC75] F. BASKETT, K.M. CHANDY, R.R. MUNTZ, J. PALACIOS
"Open, closed and mixed networks with different classes of customers" - J.ACM 22, 2, April 1975, pp. 248-260.
- [BlG79] M. BLASGEN, J. GRAY, M. MITOMA, T. PRICE
"The convoy phenomena" - ACM Operating Systems Review, 3, 2, April 1979, pp. 20-25.
- [ChS78] K.M. CHANDY, C.H. SAUER
"Approximate methods for analyzing queueing network models of computing systems" - ACM Comp. Surveys, 10, 3, Sept. 1978, pp. 281-317.
- [DeK76] P.J. DENNING, K.C. KAHN, J. LEROUDIER, D. POTIER, R. SURI
"Optimal multiprogramming" - Acta Inf. 7, 2, 1976, pp. 197-216.
- [Gil78] D.C. GILBERT
"Modelling spin locks with queueing networks" - ACM Operating Systems Review, 2, 1, Jan. 78, pp. 29-42.
- [MeP78] D. MERLE, D. POTIER, M. VERAN
"A tool for computer system performance analysis" - Proc. Int. Conf. on Performance of Computer Installations 1978, North-Holland Publishing Company, pp. 195-214.
- [MuK77] R.R. MUNTZ, G.K. KRENZ
"Concurrency in database systems - simulation study" - Proc. ACM SIGMOD Int. Conf. on Manag. of Data, Toronto, Ont., Canada, Aug. 1977.
- [Rio58] J. RIORDAN
"An introduction to combinatorial analysis" - John Wiley and Sons, Inc., New York 1958.
- [Ris77] D.R. RIES, M.R. STONEBRAKER
"Effects of locking granularity in a database management system" - ACM-TODS, 2, 3, Sept. 1977, pp. 233-246.
- [Ris79] D.R. RIES, M.R. STONEBRAKER
"Locking granule revisited" - ACM-TODS, 4, 2, June 1979, pp. 210-227.

- [Rod76] J. RODRIGUEZ-ROSELL
"Empirical data reference behaviour in database systems" - Computer 9, 11, Nov. 1976, pp. 9-13.
- [RoH75] J. RODRIGUEZ-ROSELL, D. HILDEBRAND
"A framework for evaluation of database management system" - Proc. Intern. Computing Symposium 1975, E. Gelenbe, D. Potier Eds., North-Holland Publishing Company, 1975.
- [Smi78] A.J. SMITH
"Sequentiality and prefetching in database systems" - ACM-TODS, 3, 3, Sept. 1978, pp. 223-247.
- [Ste78] W.J. STEWART
"A comparison of numerical techniques in Markov modelling" - Com. ACM, 21, 2, Feb. 78, pp. 144-151.
- [TuR76] W.G. TUEL, J. RODRIGUEZ-ROSELL
"A methodology for evaluation of database systems" - Research Report RJ 1668, IBM Research Laboratory, Yorktown Heights, New York, 1976.
- [Ver79] M. VERAN
"QNAP description language" - Research Report, Antenne Scientifique CII-Honeywell-Bull, Grenoble, 1979.

Figure 1

Figure 2

Figure 3a : $\bar{\ell}_k$ function of m , $n=4$

Figure 3b : $\bar{\ell}_k$ function of m , $n=16$

Figure 4

Figure 5 : $\mu_R(k, p)$ function of p

Figure 6

Figure 7

Figure 8

Mean number of transactions in queues

Figure 9

n	m	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8
2	2	.222	.333	.444	.555	.666	.777	.888	.999
4	2	.202	.303	.404	.505	.606	.707	.808	.909
6	2	.205	.306	.407	.508	.609	.710	.811	.912
16	2	.207	.308	.409	.510	.611	.712	.813	.914
32	2	.209	.310	.411	.512	.613	.714	.815	.916
2	4	.424	.424	.417	.403	.393	.381	.370	.359
4	4	.432	.411	.391	.369	.350	.338	.327	.316
6	4	.424	.404	.389	.369	.350	.336	.323	.310
16	4	.407	.389	.374	.357	.340	.326	.313	.300
32	4	.409	.390	.375	.358	.341	.327	.314	.301
2	6	.627	.549	.462	.356	.213	.137	.081	.037
4	6	.643	.542	.435	.301	.206	.132	.075	.034
6	6	.637	.531	.420	.300	.200	.130	.070	.030
16	6	.631	.520	.408	.286	.186	.120	.063	.027
32	6	.630	.518	.405	.283	.183	.117	.060	.026
2	16	.725	.599	.437	.294	.197	.115	.050	.027
4	16	.729	.586	.425	.282	.182	.103	.049	.024
6	16	.726	.584	.422	.280	.180	.103	.049	.024
16	16	.722	.583	.419	.277	.177	.102	.048	.023
32	16	.723	.583	.419	.277	.177	.102	.048	.023
2	32	.694	.576	.465	.331	.204	.115	.043	.020
4	32	.671	.577	.461	.329	.200	.120	.047	.022
6	32	.660	.570	.457	.323	.200	.120	.046	.021
16	32	.657	.569	.456	.322	.200	.120	.046	.021
32	32	.658	.569	.456	.322	.200	.120	.046	.021
2	64	.734	.592	.425	.271	.170	.082	.019	.012
4	64	.706	.583	.415	.269	.165	.082	.017	.009
6	64	.697	.580	.409	.264	.167	.081	.016	.009
16	64	.693	.581	.408	.262	.166	.080	.016	.009
32	64	.694	.581	.408	.262	.166	.080	.016	.009
2	128	.669	.560	.400	.250	.152	.083	.016	.007
4	128	.668	.560	.400	.250	.152	.083	.016	.007
6	128	.669	.560	.400	.250	.152	.083	.017	.007
16	128	.669	.560	.400	.250	.152	.083	.017	.007
32	128	.669	.560	.400	.250	.152	.083	.017	.007
2	256	.669	.560	.400	.250	.152	.083	.017	.007
4	256	.669	.560	.400	.250	.152	.083	.017	.007
6	256	.669	.560	.400	.250	.152	.083	.017	.007
16	256	.669	.560	.400	.250	.152	.083	.017	.007
32	256	.669	.560	.400	.250	.152	.083	.017	.007

Table I : q_k

n	m	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8
2	2	1.333	2.000	*****					
4	2	1.600	2.000	*****					
8	2	1.778	2.000	*****					
16	2	1.882	2.000	*****					
32	2	1.939	2.000	*****					
2	4	1.600	3.000	3.833	4.000	*****			
4	4	2.286	3.671	3.969	4.000	*****			
8	4	2.909	3.910	3.999	4.000	*****			
16	4	3.369	3.977	3.999	4.000	*****			
32	4	3.657	3.996	4.000	4.000	*****			
2	8	1.778	3.800	5.132	6.596	7.897	7.999	7.977	8.000
4	8	2.909	5.613	7.196	7.950	7.969	7.993	7.999	8.000
8	8	4.267	7.037	7.985	7.986	7.998	8.003	8.000	8.000
16	8	5.865	7.740	7.995	7.999	8.000	8.000	8.000	8.000
32	8	6.864	7.963	7.999	8.000	8.000	8.000	8.000	8.000
2	16	1.909	3.780	5.899	7.627	9.213	10.954	12.616	13.120
4	16	3.369	6.899	9.626	12.327	14.562	15.896	15.997	15.966
8	16	5.865	10.361	13.864	15.809	15.939	15.990	15.997	15.999
16	16	7.256	12.664	15.727	15.879	15.999	16.000	16.000	16.000
32	16	10.096	15.370	15.976	15.999	16.000	16.000	16.000	16.000
2	32	1.939	3.975	5.916	7.773	9.654	11.568	13.473	15.372
4	32	3.657	7.274	11.061	14.344	17.761	21.061	24.192	27.163
8	32	6.864	12.036	18.690	23.914	28.140	31.052	31.707	31.987
16	32	10.096	20.235	27.220	33.987	39.900	41.999	41.999	41.999
32	32	16.254	30.994	41.760	51.971	59.990	62.999	62.999	62.999
2	64	1.969	3.937	5.935	7.871	9.826	11.803	13.761	15.717
4	64	3.821	7.671	11.620	15.617	19.691	23.733	27.663	31.463
8	64	7.211	14.737	21.345	28.716	36.938	44.739	51.480	57.392
16	64	12.962	25.336	36.040	47.137	58.356	69.000	77.862	84.987
32	64	21.840	43.039	63.964	81.685	97.836	111.999	123.999	134.999
2	128	1.996	3.969	5.953	7.976	9.993	11.993	13.996	15.996
4	128	3.969	7.913	11.717	15.616	19.512	23.404	27.297	31.187
8	128	7.895	15.146	22.670	30.103	37.645	45.177	52.667	60.100
16	128	15.722	30.286	42.761	55.205	67.618	80.000	92.317	104.597
32	128	25.761	50.313	73.103	95.633	117.233	138.670	159.999	179.999
2	256	1.992	3.991	5.975	7.967	9.967	11.957	13.960	15.957
4	256	3.984	7.907	11.750	15.600	19.451	23.311	27.163	31.000
8	256	7.967	15.860	23.741	31.560	39.366	47.166	54.957	62.700
16	256	15.914	31.176	45.163	61.117	76.170	91.784	106.87	121.999
32	256	25.844	50.708	75.397	101.800	127.265	152.190	177.160	201.200

Table II : $\bar{\ell}_k$


```

/DECLAR/QUEUE CPU,DU1,DU2,DU3;
CLASS ACTIVE,REQUEST;
REAL XA,XR,Y,NA,NR,TA(100),TR(100),M,N;
INTEGER KTR,IA,IR,I;

&
/STATION/NAME=CPU;
SCHED=PS;
SERVICE(ACTIVE)=EXP(XA);
SERVICE(REQUEST)=EXP(XR);
TRANS=DU1,1,DU2,1,DU3,1;
INIT(ACTIVE)=IA;
INIT(REQUEST)=IR;

&
/STATION/NAME=DU1;
SERVICE=EXP(Y);
TRANS=CPU;

&
/STATION/NAME=DU2;COPY=DU1;

&
/STATION/NAME=DU3;COPY=DU1;

&
&
/PARAM/CLASS=CPU;
/EXEC/BEGIN
XA:=0.05;
XR:=0.01;
N:=4;
M:=1;

&
NA:=4;
NR:=M*N/(M+N-1);

&
Y:=0.05;
KTR:=5;

&
IA:=0;
WHILE IA<=KTR DO BEGIN
  IR:=0;
  WHILE IR<=KTR DO BEGIN
    ANALYTIC;
    I:=(KTR+1)*IA+IR+1;
    TA(I):=NA/MTHRUPUT(CPU,ACTIVE);
    TR(I):=NR/MTHRUPUT(CPU,REQUEST);
    PRINT(I,IA,IR,TA(I),TR(I));
    IR:=IR+1;
  END;
  IA:=IA+1;
END;

END;

**** ANALYTIC RESOLUTION ****

****SUBCHAIN 1 TOTAL NUMBER OF CUSTOMERS= 3
****SUBCHAIN 2 TOTAL NUMBER OF CUSTOMERS= 5
*****
* NAME * SERVICE * BUSY PCT * CUST NB * RESPONSE * THRUPT *
*****
* CPU * .2027E-01* .8535 * 2.637 * .6262E-01* 42.12 *
* (ACTI) * .5000E-01* .5404 * 1.528 * .1413 * 10.81 *
* (REQU) * .1000E-01* .3131 * 1.110 * .3545E-01* 31.31 *
* * * * *
* DU1 * .5000E-01* .7020 * 1.788 * .1273 * 14.04 *
* * * * *
* DU2 * .5000E-01* .7020 * 1.788 * .1273 * 14.04 *
* * * * *
* DU3 * .5000E-01* .7020 * 1.788 * .1273 * 14.04 *
* * * * *
*****
SPACE USED : 3/100
**** END OF ANALYTIC RESOLUTION ****

```

Table III

```

/DECLAR/QUEUE TERMINAL,RESREQ,RESACT,BLOCK;INTEGER
      K1,K2;REAL TREF,TREQUEST,TACTIVE,QQ;
      REAL Q(10)=1.0,0.94,0.883,0.828,0.776,0.727,0.68,0.635
              ,0.592,0.332;
&
$ MACRO IND(I)
      IA:=CUSTNB(RESACT);
      IR:=CUSTNB(RESREQ);
      I:=(KTR+1)*IA+IR+1;
$ END
&
/STATION/NAME=TERMINAL;
      TYPE=DELAY;
      SERVICE=EXP(TREF);
      TRANS=RESREQ(REQUEST);
      INIT(ACTIVE)=KTR;
/STATION/NAME=RESREQ;
      SERVICE=BEGIN
              I:=CUSTNB(RESACT);
              QQ:=Q(I+1);
              $IND(I)
              EXP(TR(I));
      END;
      TRANSIT=RESACT(ACTIVE),QQ,BLOCK;
/STATION/NAME=RESACT;
      SERVICE=BEGIN
              K1:=1;
              $IND(I)
              EXP(TA(I));
      END;
      TRANSIT=TERMINAL;
/PARAM/OPTION=RESULT;TEST=BEGIN
      IF K1=1 THEN BEGIN K1:=0;MOVE(BLOCK,RESREQ);END;
      END;
/EXEC/BEGIN
      NETWORK(TERMINAL,RESREQ,RESACT,BLOCK);
      K1:=0;
      TREF:=1.5;
      MARKOV;
      PRINT:(CUSTNB(BLOCK),MTHRUPUT(BLOCK),MRESPONSE(BLOCK));
      END;

```

Table IV