

On Merging Partitioned Databases

David D. Wright

Department of Computer Science
Cornell University
Ithaca, New York 14853

Abstract

Partitioning of a distributed data base requires either that update activity be restricted or that a strategy for conflict resolution and partition merging be used once communication is restored. The graph-theoretic approach used by Davidson follows the latter approach and can be used to show that finding an optimum solution to the general problem is NP-complete. We give several methods of reducing the size of the graphs involved. Two open subproblems are shown to be NP-complete, while an extension of a known polynomial-time subproblem is given. Simulation results are used to study both the amount of compression achieved by the graph reduction techniques and their effects on heuristics for the problem. In addition, some modifications are made to existing heuristics to improve their performance. A simple probabilistic model is developed and compared to the simulation results.

1. Assumptions and Definitions

A distributed database (DDB) consists of a set of *data-items* $\{d_1, \dots, d_M\}$ and a set of *sites* $\{s_1, \dots, s_N\}$ that are connected by bidirectional communication links. We assume that during normal operation any site can send a message to any other site. The values of data-items are manipulated by *transactions*, which have the property that their actions are atomic from the point of

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

view of any other transaction. Every transaction T has associated with it two sets, $READSET(T)$ and $WRITESET(T)$, which are, respectively, the sets of items read and written by T . We assume that $WRITESET(T) \subseteq READSET(T)$, and that data are fully replicated, i.e. that there is a copy of every data-item at every site.

A *partition* of the DDB is a maximal subset of communicating sites. Site or communication link failures may separate the DDB into more than one partition, whereas site or link recoveries may require that partitions be merged. The system is said to be *partitioned* at any time that it is composed of more than one partition. We assume that within a partition the transaction activity corresponds to some serial schedule, and that this schedule can be derived by any member of the partition (using the history described below).

1.1. Description of the Problem

Strategies to allow a partitioned DDB to continue functioning generally fall into one of two categories. Most previous methods have been deterministic, guaranteeing that conflicts will be avoided. This category includes voting schemes [Giff79, Thom78], token passing, and so on. A problem with these schemes is that failures may occur in such a way that no partition can perform updates; for example, in a voting scheme, no partition may have a majority. In contrast, the second category of algorithms is probabilistic and allows each partition to perform updates that might conflict with updates in another partition. Such strategies are called *optimistic* by Davidson [Davi82] because they assume that the number of conflicts will be small. When communication between two or more partitions is restored, the partitions must compare their actions during partitioning, detect

conflicts, then back out (undo) transactions until no conflicts remain. To simplify the discussion, we shall assume that only two partitions are to be reconnected at a time. However, the method described generalizes to any number of partitions.

The approach used here requires that whenever the DDB is partitioned, each partition P_i must maintain a history of the READSET and WRITESET of all transactions T_{ij} that have committed (successfully completed) in that partition since it was formed, along with the order in which the transactions committed. When two partitions are reconnected, each partition derives a serial schedule H_i for its partition using the history it has maintained (we assumed this to be possible earlier). The serial schedules are then used to construct a precedence graph.

Definition 1.1 [Davi82]. Let $H_1 = T_{11}, T_{12}, \dots, T_{1n}$ and $H_2 = T_{21}, T_{22}, \dots, T_{2m}$ be serial schedules from partitions P_1 and P_2 respectively. The *precedence graph* $G(H_1, H_2) = (V, E)$ is the digraph defined by:

- (1) $V = \{T_{11}, \dots, T_{1n}\} \cup \{T_{21}, \dots, T_{2m}\}$.
- (2) $E = \{\text{Dependency Edges}\} \cup \{\text{Precedence Edges}\} \cup \{\text{Interference Edges}\}$, where
 - (a) **Dependency Edges**¹—there is a dependency edge $T_{ai} \rightarrow T_{ak}$ iff $i < k$ and there is a data-item d such that $d \in \text{WRITESET}(T_{ai}) \cap \text{READSET}(T_{ak})$ and $(\forall j: i < j < k: d \notin \text{WRITESET}(T_{aj}))$.
 - (b) **Precedence Edges**—there is a precedence edge $T_{ai} \rightarrow T_{ak}$ iff $i < k$, there is no dependency edge $T_{ai} \rightarrow T_{ak}$, and there is a data-item d such that $d \in \text{READSET}(T_{ai}) \cap \text{WRITESET}(T_{ak})$ and $(\forall j: i < j < k: d \notin \text{WRITESET}(T_{aj}))$.
 - (c) **Interference Edges**—there is an interference edge $T_{1i} \rightarrow T_{2j}$ iff there is a data-item d such that $d \in \text{READSET}(T_{1i}) \cap \text{WRITESET}(T_{2j})$. (Similarly for $T_{2i} \rightarrow T_{1j}$.)

For each vertex $v \in V$, the *dependency set* of v , $\text{dep}(v)$, is the set of all vertices v' in V s.t. there is a path of dependency edges from v to v' . For example, in Fig. 1.1, the dependency set of c is $\{a, d\}$, and $\text{dep}(c) = \{a, b, c, d\}$. If v is backed out, then $\text{dep}(v)$ must be backed out, since the values read by the transactions in $\text{dep}(v)$ depend on the values written by v . We will use

¹Called *ripple edges* in [Davi82].



Fig. 1.1

the convention that if S is a set of vertices, $\text{dep}(S)$ is the union of the dependency sets of the members of S .

[Davi82] contains a proof that $G(H_1, H_2)$ is acyclic iff the combined set of transactions is serializable. Thus, if G contains cycles and we wish to find a serializable subset of $T_1 \cup T_2$, some transactions (and their dependency sets) must be backed out until the graph is acyclic. A set S of vertices with the property that its removal leaves the graph acyclic is called a *feedback vertex set*.

We can now state the *Partition Merging Problem* (PM): given a precedence graph $G = (V, E)$ and an integer $K > 0$, is there a set $S \subseteq V$ such that $\text{dep}(S)$ is a feedback vertex set of G and $|\text{dep}(S)| \leq K$? Such a set S is called a *transaction backout set*. The quantity $|\text{dep}(S)|$ is called the *weight* of the set S .

2. Some Properties of Precedence Graphs

This section contains some important lemmas about the properties of precedence graphs. The point of these lemmas is that certain edges may be added or deleted without affecting the solution to an instance of PM.

Lemma 2.1: Suppose $G = (V, E)$ is a precedence graph, $a, b, c \in V$, $(a, b) \in E$ is any type of edge, $(b, c) \in E$ is a dependency edge, and $(a, c) \notin E$. Let $G' = (V, E \cup \{(a, c)\})$, where (a, c) is of the same type as (a, b) . Then $S \subseteq V$ is a backout set of $G \Leftrightarrow S$ is a backout set of G' .

Proof: (\Leftarrow) Trivial, since G is a subgraph of G' . (\Rightarrow) Suppose S is a backout set of G . Any cycle that includes (a, c) has a corresponding cycle in which (a, c) is replaced by (a, b) and (b, c) . All of the latter are broken

by S . This could be done by removing a , b , c , or some other nodes on the cycles involving (a,b) and (b,c) . However, any of those removals will also break all cycles involving (a,c) . Thus, S must be a backout set of G' . \square

One useful consequence of this lemma is that an interference edge need only be drawn from a transaction T to the first transaction in the other partition that writes any item read by T . This simplifies constructing precedence graphs.

Employing the above lemma will never increase the number of strongly-connected components in the graph. In particular, individual nodes will never become disconnected. The next lemmas illustrate that in some cases, edges can be removed in such a way that the number of components may increase.

Lemma 2.2: Let $G=(V,E)$ be a precedence graph, where $(a,b)\in E$ and $(b,a)\in E$ are interference edges, $c\in dep(a)$, $d\in dep(b)$, $(c,d)\in E$, and $(c,d)\neq(a,b)$. Then $S\subseteq V$ is a backout set of $G \Leftrightarrow S$ is a backout set of $G'=(V,E-\{(c,d)\})$.

Proof: (\Rightarrow) Trivial, since G' is a subgraph of G . (\Leftarrow) Suppose S is a backout set of G' but not of G . Any cycle not broken in G must include (c,d) , since all other edges of G are in G' . But the 2-cycle² involving a and b is in G' and can only be broken by deleting a or b . Deleting either of these will result in removing (c,d) , so that edge cannot be involved in a cycle in G , a contradiction. Hence, S must be a backout set of G . \square

Lemma 2.3: Let $G=(V,E)$ be a precedence graph, where $(a,b)\in E$ and $(c,b)\in E$ are interference edges, $c\in dep(a)$, $c\neq b$. Suppose that all cycles involving (c,b) pass through a . Then $S\subseteq V$ is a backout set of $G \Leftrightarrow S$ is a backout set of $G'=(V,E-\{(c,b)\})$.

Proof: (\Rightarrow) Trivial, since G' is a subgraph of G . (\Leftarrow) Suppose that S is a backout set of G' . Note that every cycle C involving (c,b) has a corresponding cycle in which the section of C from a to c and (c,b) is replaced by (a,b) . This latter cycle is broken by S . This could be done by removing a , b , or some other transaction on the path from b to a ; however, any of these removals will also break C . \square

One easy method of constructing an equivalent precedence graph is to remove all edges (and thus vertices)

that do not lie on a cycle. This can be done by finding the strongly connected components of the graph and deleting all edges that do not connect members of the same component. This requires node weights (if needed) to be stored at the nodes, as there would not necessarily be any way to compute the weights from the reduced graph.

3. The Complexity of Partition Merging

This section contains several theoretical results about the partition merging problem. The first is

Theorem 3.1: The partition merging problem is NP-complete.³

Proof: We will show that an instance of the feedback vertex set problem (does an arbitrary digraph G have a feedback vertex set of size $\leq K$?) is polynomially transformable to an instance of the partition merging problem. (For a proof that the feedback vertex set problem is NP-complete, see e.g. [AHU74].) Let $G=(V,E)$ be an arbitrary directed graph. It is obvious that we can determine in polynomial time if some set $S\subseteq V$ is a transaction backout set of G . We will show how to construct in polynomial time a graph $G'=(V',E')$ such that G' is a precedence graph with a transaction backout set of weight $\leq K$ iff G has a feedback vertex set of size $\leq K$.

Define G' as follows: let $V' = V \cup \{v_{ab} : (a,b)\in E\}$, and $E' = \{(a,v_{ab}), (v_{ab},b) : (a,b)\in E\}$. That is, replace each edge $(a,b)\in E$ by two edges, (a,v_{ab}) and (v_{ab},b) . (See Fig. 3.1.)

Let $T_1=V$ and $T_2=(V'-V)$, with H_1 and H_2 arbitrary orderings of T_1 and T_2 . Let E' consist solely of interference edges. Then G' is a precedence graph, and $(\forall v\in V' : dep(v)=\{v\})$.

Note that the cycles of G and G' are in a 1-to-1 correspondence. Clearly, any feedback vertex set S of G is a feedback vertex set (and thus a transaction backout set) of G' , and since each $v\in V'$ has a dependency set consisting only of itself, $|S|\leq K$ iff $weight(S)\leq K$.

Now suppose that S is a transaction backout set of G' . Then there is a transaction backout set that is no larger than S and contains only members of V . To see this, notice that if there were some $v\in S$ with $v=v_{ab}$, we

²A k -cycle is a cycle of length k .

³[Davi82] contains a similar, independently developed, proof.

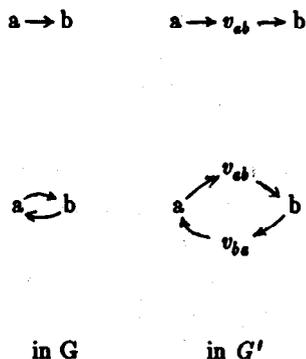


Fig. 3.1

could replace it in S by b , since any cycle that includes v_{ab} must include b . Call S' the set derived from S by doing all such replacements. Since $(\forall v \in V' : \text{dep}(v) = \{v\})$, $|S'| \leq K$ and S' is a feedback vertex set of G . \square

This proof holds only if the number of items in the database is unbounded, since the number of items needed is of the order of the number of edges in the graph G . If the number of items is a constant, the problem can be solved in polynomial time.⁴ [WrSk83]. However, it seems likely that in those environments for which the optimistic protocol is suited, the number of data items will be much larger than the number of transactions.

In a real database system, one partition might be "privileged" in that its transactions would never be backed out. This restriction does not help, however; since the construction used in the previous proof only removes nodes from one partition, we can conclude

Corollary 3.1: The partition merging problem with one partition ineligible for backout is NP-complete.

Even if a problem is NP-complete, it may have subproblems that can be solved in polynomial time. [Davi82] contains one such example:

Theorem 3.2: Let $G=(V,E)$ be a precedence graph for two partitions such that all edges are either dependency edges or are interference edges that lie on 2-cycles. Then a solution to PM can be found in polynomial time. \square

Davidson's proof actually allows a more general conclusion. Suppose $P \subseteq E$ be the set of precedence edges in E . Then any $S \subseteq E$ is a backout set for G iff it is a

⁴Result due to M. Krentel of Cornell University.

backout set for $G'=(V,E-P)$. To see this, suppose S is a backout set of G' . Then it must break all 2-cycles in G' (all of which are also in G). The only way to do this is to remove one or both vertices lying on each 2-cycle. This has the effect of removing all interference edges in G' (hence G). Since a graph with no interference edges can have no cycles, S is a backout set of G .

The above theorem is particularly useful since, for most environments, we would expect that very few cycles will be left after all 2-cycles are broken.

Unfortunately, Theorem 3.2 does not extend to more than 2 partitions.⁵ We will show this using a restricted form of the problem "does G have a vertex cover⁶ of size $\leq K$?" (VC), which is NP-complete, even for cubic planar graphs⁷ [GaJo79]. We will first show that a restricted version of VC on cubic planar graphs is NP-complete. Next, we will show how to convert an instance of this problem to an instance of PM that has only three partitions and all interference edges lying on 2-cycles.

Lemma 3.3: The vertex cover problem is NP-complete for cubic planar graphs that do not have K_4 as a component.

Proof: Let G be a cubic planar graph. We can easily find the K_4 components of G in polynomial time. (They are simply those components containing exactly 4 vertices.) Assume that there are m such components, and let G' be the subgraph of G formed by deleting all those components. Since K_4 has a vertex cover consisting of any 3 of its vertices, G' has a vertex cover of size K iff G has a cover of size $K+3m$. \square

Lemma 3.4: Every cubic planar graph $G=(V,E)$ that does not have K_4 as a component can be 3-colored in polynomial time.

Proof: see [Lova75].

Theorem 3.5: PM is NP-complete for precedence graphs of 3 or more partitions even if all interference edges lie on 2-cycles.

⁵It also does not extend to precedence graphs that contain no 2-cycles and all of whose edges lie on 3-cycles [WrSk83].

⁶A subset S of V such that every edge in E is incident on some member of S .

⁷A graph is cubic if all of its vertices have degree 3.

Proof: Let $G=(V,E)$ be a cubic planar graph that does not contain K_4 . We will construct in polynomial time a precedence graph $G'=(V,E')$ for 3 partitions with the property that G has a vertex cover of size $\leq K$ iff G' has a backout set of size $\leq K$. Define G' as follows: replace each edge $\{a,b\} \in E$ in G with two directed edges (a,b) and (b,a) in E' . Let each edge in E' be an interference edge. Construct a 3-coloring for G , and use the colors as names for the partitions in G' . Then there will be no interference edges connecting two members of the same partition. Suppose $S \subseteq V$ is a vertex cover of G but not a backout set of G' . But every edge in G (hence G') is incident on a member of S and hence no cycle can exist. Contrariwise, suppose $S \subseteq V$ is a backout set of G' . In particular, this means that all 2-cycles of G' are broken by S . But every edge in G' is on a 2-cycle (corresponding to an edge in G) and hence S must be a vertex cover of G . \square

4. Simulation of Backout Strategies

The fact that a problem is NP-complete does not mean that instances of it difficult to solve, only that they may be difficult to solve optimally. For the optimistic protocol to be of practical importance, however, efficient heuristics that usually give good results must be found.

For some NP-complete problems, efficient polynomial-time heuristics with good worst-case error bounds have been found. (Examples can be found in [GaJo79], Chapter 6; errors are measured as the relative size of the solution found by the heuristic to the actual minimum.) For PM, straightforward methods such as iteratively backing out the transaction of lowest weight until the graph is acyclic have worst-case errors of $O(n)$ for a graph of n nodes. In fact, even the best of the known heuristics has an error bound of $O(n/\log(n))$ [WrSk83], which is too large to be of any practical significance. Thus, the decision was made to restrict attention to fast heuristics with $O(n)$ possible error to see if their performance was adequate.

To gain understanding of heuristic solutions, and to compare the relative performance of different solutions, computer simulation of random transactions in a 2-partition DDB was used. (This approach was also taken in [Davi82], which contains independent results of simulating several strategies on 2-partition DDBs. The results below reflect a more extensive set of observations.) Details on the simulation are contained in the Appendix.

4.1. Using Reduced Graphs

One important simulation goal was to discover what effect the use of reduced graphs would have on the backout rates of various heuristics. It seemed possible that the graph reductions might either cause improvement (by eliminating spurious nodes from consideration) or degradation (by eliminating significant nodes from consideration). In practice, neither of these occurred: results with extensive reduction (making full use of the graph reduction lemmas) and simple reduction (strongly-connected component construction only) had no significant differences between them.

Although reduction does not cause major improvements in backout rates, it has other merits, as the running time of some heuristics can be significantly improved by applying them to reduced graphs. Davidson's method for optimally breaking 2-cycles has time complexity $O(n^3)$ for an n -vertex graph. The advantages of reduction here are obvious.⁸ In addition, construction of precedence graphs was simplified by Lemma 2.1, since interference edges need only be drawn to the first node that writes a given item.

A number of simulation runs were made to estimate the amount of compression that reduction can give. The results were very encouraging; even under high-backout conditions (backout rates around 30%), the number of nodes in the graph was typically reduced by more than 40%. Under low-backout conditions (rates of 10% or less), the number was reduced far more; 90% was typical. About 3/4 of the compression can be attributed to strongly-connected component construction, with the rest due to redundant interference edge elimination.

Preliminary simulation results suggested that some heuristics would perform better if they were restricted to backing out transactions in only one partition. Since restriction is not symmetric (restriction to one partition does not necessarily give the same results as restriction to the other), the problem is then to pick the better partition. Graph reduction provides a partial solution: even if the size of each partition is initially the same, in general the sizes differ after reduction. The smaller partition can then be used by heuristics that only consider one partition.

⁸Some restrictions on reduction would be necessary in this case.

The amount of improvement given by this modification over the original strategies varied depending on the backout rate. When the rate was low ($< 5\%$), differences were insignificant. However, as the rate rose, the modified strategies began to gain; a typical example had the backout rate dropping from 28% to 24%⁹, an improvement of 14%. This is particularly significant because the improvement is greater than the differences between the individual strategies, which were generally less than 5%.

Unfortunately, relative partition size is only a guide. If the partitions were approximately the same size after reduction, applying the heuristics to the larger partition sometimes gave better results. In cases like this, trying the heuristics on both partitions and taking the better result improved performance by about 10% in high-backout situations.

4.2. Parameter Sensitivity

A change in some simulation parameters affects the output more than a change in others. The parameters for this simulation were:

- (1) $N1, N2$ number of transactions in partition 1 and 2, respectively
- (2) M data-items in the database
- (3) I average no. of data-items referenced by a transaction
- (4) RO average percentage of read-only transactions
- (5) U average percentage of data-items updated in each non-read-only transaction

The simulations provided the following results. (The Appendix contains graphs illustrating some of these.)

- o Holding the ratio $(N1 + N2)/M$ constant and all other parameters fixed (including the ratio $N1/N2$) gives a constant backout rate, as was expected.
- o The backout rate is sensitive to changes in RO throughout its range (see Appendix).

⁹In this and most other cases, the number of simulation samples was 100 to 150, which was sufficient to make the 95% confidence intervals on the backout rates less than one percentage point.

- o Changes in U cause approximately linear change in the backout rate from $U=0\%$ to about $U=50\%$; thereafter, change increases the rate only slightly (see Appendix).
- o The backout rate is inversely proportional to M for rates below about 30%.
- o The probability distribution of data-item references has a relatively small on the backout rate. Even an 80/20 distribution (20% of the items referenced 80% of the time) increases the backout rate by only a few percentage points over a uniform distribution. Thus, the use of a uniform distribution by the simulations is reasonable.
- o If WRITESETs between partitions do not overlap, backout rates drop dramatically (see Appendix).

4.3. Using Optimistic Commit

For optimistic commit to be worthwhile in a DDB, the costs of the protocol must be outweighed by the costs of delaying transactions. We can express this relationship for n transactions as follows:

$$D \frac{n}{2} > A(n) + (B + C)b(n)n, \text{ where}$$

$A(n)$ is the cost of running the backout selection algorithm

B is the average cost of backing out a transaction (may include a penalty cost)

C is the average cost of re-executing a transaction

D is the average cost of delaying a transaction

$b(n)$ is the fraction of transactions backed out

The $\frac{n}{2}$ term on the left side of the inequality reflects that an ordinary DDB, when partitioned, would typically be able to run about half of the transactions submitted (using voting or some other method). In the strategies used in our simulations, $A(n) = An$. In this case, we can rearrange this equation to $\frac{D-2A}{2(B+C)} > b(n)$. For values of n satisfying this inequality, it will be advantageous to use optimistic commit. Most of the parameters involved should be fairly easy to estimate for a given database system.

4.4. A Probabilistic Model

In this section we develop a probabilistic model that can be used to estimate the backout rate of a set of transactions as a function of the parameters used by the computer simulation. We consider two types of conflict – 2-cycles caused by write-write interference and direct dependency effects of those 2-cycles – and estimate the number of nodes removed due to both types.¹⁰

First, we estimate NTC , the number of nodes in one partition that are in 2-cycles. (We will consider $P1$ for purposes of exposition. In addition, we will use RO and U as fractions rather than percentages in order to simplify the formulas.) Only an update transaction, of which there are $N1(1-RO)=NRW$, can be involved in a 2-cycle. The average number of items updated by such a transaction is $I*U$. As long as the total number of update references by transactions in $P2$, $N2*I*(1-RO)*U=NUP$, is $\ll M$, we can use NUP as the number of items updated in $P2$. Then for an update transaction T in $P1$, the probability that a given item T writes is not written in $P2$ is $\left(\frac{M-NUP}{M}\right)$, and the probability that no item T writes is written in $P2$ is $\left(\frac{M-NUP}{M}\right)^{I*U}$. Thus, the probability that some item written by T is also written in $P2$ (i.e. that T is in a 2-cycle) is $1-\left(\frac{M-NUP}{M}\right)^{I*U}=PTC$. The expected number of transactions in $P1$ involved in 2-cycles is simply $PTC*NRW=NTC$.

Second, we estimate NRD , the number of nodes backed out because they are dependency children of nodes in NTC . The relevant nodes are those that are not involved in 2-cycles, but have dependency parents that are. Let XTC be the number of nodes not involved in 2-cycles, i.e. $N1-NTC$. The probability that a node in $P1$ is not involved in a 2-cycle is $\frac{XTC}{N1}$. Only update transactions can have dependency descendants or be involved in 2-cycles, and the expected number of update transactions preceding T_k , the k th node in the serial schedule for $P1$, is $(k-1)(1-RO)$. The probability that a given predecessor of T_k is both in a 2-cycle and a dependency parent of T_k is $\frac{NTC}{N1}\left(1-\left(\frac{M-I}{M}\right)^{I*U}\right)$, so the proba-

bility that a predecessor is either not in a 2-cycle or not a dependency parent of T_k is $1-\frac{NTC}{N1}\left(1-\left(\frac{M-I}{M}\right)^{I*U}\right)$, and the probability that all predecessors are either not in a 2-cycle or not dependency parents of T_k is $\left(1-\frac{NTC}{N1}\left(1-\left(\frac{M-I}{M}\right)^{I*U}\right)\right)^{(k-1)(1-RO)}$. The probability that T_k is counted in NRD is then

$$(*) \ 1-\left(1-\frac{NTC}{N1}\left(1-\left(\frac{M-I}{M}\right)^{I*U}\right)\right)^{(k-1)(1-RO)}$$

In order to find the expected number of nodes removed strictly due to dependency effects, we must sum $(*)$ from $k=1$ to $k=N1$ and multiply by the expected fraction of non-2-cycle nodes $(XTC/N1)$. First, we define $\alpha=\left(1-\frac{NTC}{N1}\left(1-\left(\frac{M-I}{M}\right)^{I*U}\right)\right)^{(1-RO)}$. Then the expected number of nodes removed due to dependency effects is $\frac{XTC}{N1}\sum_{k=1}^{N1}1-\alpha^{k-1}$, or $\frac{XTC}{N1}\left(N1-\sum_{k=1}^{N1}\alpha^{k-1}\right)$. As long as the problem is non-trivial (i.e. $U > 0$, $RO < 100\%$, etc.) $0 < \alpha < 1$, and we can use the identity $\sum_{k=1}^{\infty}\alpha^{k-1}=\frac{1}{1-\alpha}$ to rewrite the summation as $\left(\frac{XTC}{N1}\right)\left(N1-\frac{1-\alpha^{N1}}{1-\alpha}\right)=NRD$, making the total number of nodes backed out $NTC+NRD$.

4.5. Comparison with Simulation

The probabilistic model was compared with a wide variety of simulation runs in order to determine the agreement between the two. Agreement was excellent, generally within 10%. Since the differences in practice would be affected at least as much by factors such as the variance of the transaction size, this is quite satisfactory.

5. Discussion

The significant results of this paper are

- (1) New techniques allowing significant reduction in the size of precedence graphs.
- (2) New complexity results, particularly the case of multiple partitions in which all interference edges lie on 2-cycles.
- (3) Improvements in heuristics and analysis of the impact that reduction has on heuristic performance.

¹⁰[Davis2] considers the same effects, but using a much more restricted model. A more detailed model is contained in [WrSk83].

The graph reduction results are important for several reasons. First, they simplify construction of the precedence graph, since interference edges need only run to the first node that writes a given item, not to all of them. This reduces the memory needed to construct the graph, since the graph-construction program needs only a table of "first writes," not a list of all writes. Second, they make the graph smaller. Third, their use does not degrade the performance of the heuristics studied in this paper. Fourth, and perhaps most important, restricting heuristics to one partition only and using the smaller partition after reduction can improve the backout rate of the heuristics studied.

6. Acknowledgements

I am particularly grateful to Dale Skeen for his many valuable contributions to this work. I would also like to thank M. Krentel, J. Gilbert, and B. Aspvall for their helpful advice, and K. Birman for reading an earlier version of this paper. Several useful suggestions were also made by the referees. This work is supported in part by NSF Grant MCS-8103605.

7. References

- [AHU74] Aho, A.V., J.E. Hopcroft, and J.D. Ullman, *The Design and Analysis of Computer Algorithms*, Addison-Wesley, 1974.
- [Davi82] Davidson, S., "An Optimistic Protocol for Partitioned Distributed Database Systems", PhD Thesis, Princeton University, Dept. of EECS, 1982.
- [Giff79] Gifford, D.K., "Weighted Voting for Replicated Data", 7th Symposium on Operating System Principles, Dec. 1979, pp.150-162.
- [GaJo79] Garey, M.R., and D.S. Johnson, *Computers and Intractability*, W.H. Freeman & Co., 1979.
- [Lova75] Lovasz, L., "Three Short Proofs in Graph Theory", *J. Combinatorial Theory B*, 19, 111-113.
- [Thom78] Thomas, R.H., "A Solution to the Concurrency Control Problem for Multiple Copy Databases," *IEEE Compcon*, Spring 1978, pp.56-62.
- [WrSk83] Wright, D.D., and Skeen, D. "Managing Partitioned Databases," Technical Report TR 83-547, Dept. of Computer Science, Cornell University, March 1983.

Appendix

This section discusses some details of the simulation and provides sample graphs that show the sensitivity of the simulation output to changes in the parameters. For a more complete discussion, see [WrSk83].

The simulations used pseudo-randomly generated transactions with the parameters defined earlier. Data-item references were uniformly distributed over the database in the simulation.¹ Update-set sizes had a binomial distribution and transaction sizes were exponentially distributed.

A wide variety of strategies were used in the simulations, but two emerged as preferable. Both were originally described in [Davi82]. They are:

- (a) Davidson strategy #2. Iteratively break all 2-cycles by removing the node with lower weight.
- (b) Davidson strategy #3. Break all 2-cycles by removing the node in partition 1. (As described in an earlier section, the performance of this heuristic can be improved by applying it to the smaller partition.)

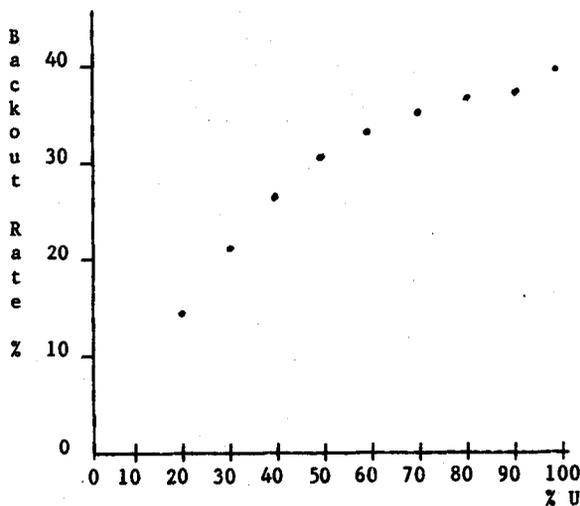
If the precedence graph was not acyclic after all 2-cycles had been broken, a minimum-weight node removal strategy was used to finish the job.

All other strategies considered had performance essentially similar to Davidson #3, although they were slightly inferior to it. While Davidson #3 did not always give the best performance, it was typically as good or better than any of these other strategies on a given example.

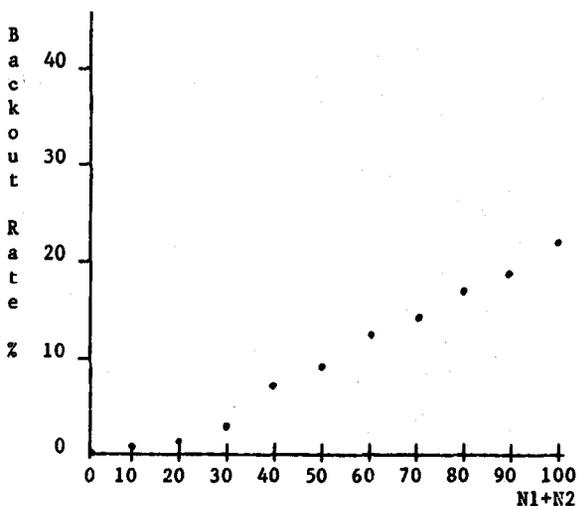
The Davidson #2 strategy was distinctly different. It had better performance under low-backout conditions, but was worse under high-backout conditions. The point at which it became worse varied, but was usually around a backout rate of 15-20%. Since optimistic commit probably would be used under low-backout conditions, this strategy would be the one of choice. (There remains the possibility of using Davidson's optimal 2-cycle breaking algorithm. However, since this has time complexity $O(n^3)$, it would be infeasible for large values of n . A faster approach would be to try several heuristics and use the best result.)

¹Except in experiments measuring the effects of changing this distribution.

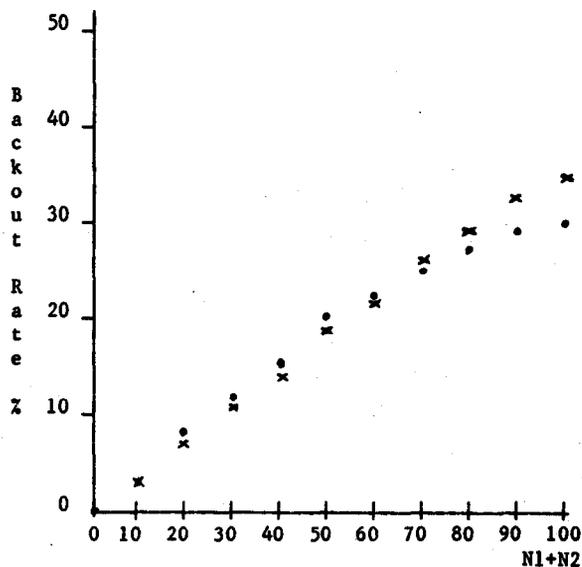
One note about costs should be made. The backout rates in this report are simply the number of transactions backed out as a percentage of all transactions. This is not an accurate reflection of all environments. In particular, since backing out a readonly transaction does not require changes to the database, it may be more reasonable to give readonly transactions a weight of zero. Simulations run using this cost function showed relatively little change from the regular costs when backout rates were low. However, as the rates rose, the modified function had increasingly better performance. For example, when the original rate was 8%, the modified cost was 6%. On the other hand, an original rate of 30% dropped to 17%. These results reflect the number of readonly transactions that are in the dependency sets of backed-out update transactions.



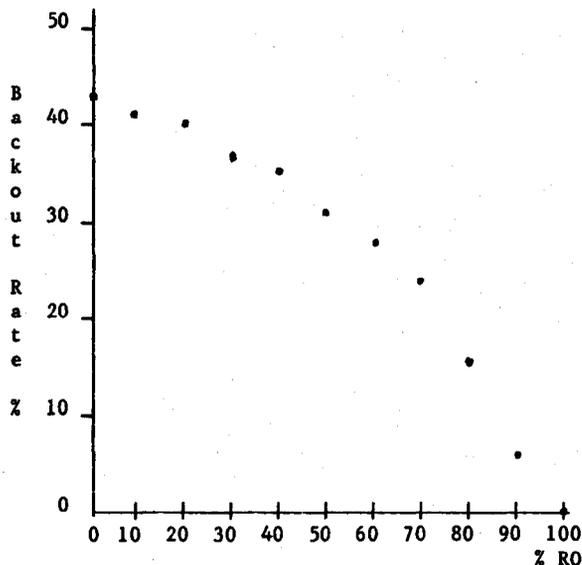
Sensitivity of Backout Rate to U
 $N1 = N2 = 50, M = 100, RO = 50\%, I = 5$
 Strategy: Davidson #3



Backout Rate for Non-overlapping Writesets
 (each Writeset was 50% of database)
 $M = 100, U = 50\%, RO = 50\%, I = 5, N1/N2 = 1$
 Strategy: Davidson #3



Sample Backout Curves
 $M = 100, RO = 50\%, U = 50\%, I = 5, N1/N2 = 1$
 Strategies: x -- Dav. #2 o -- Dav. #3



Sensitivity of Backout Rate to RO
 $N1 = N2 = 50, U = 50\%, M = 100, I = 5$
 Strategy: Davidson #3