

# Dataflow Query Execution in a Parallel Main-Memory Environment

Annita N. Wilschut

Peter M.G. Apers

University of Twente

P.O.Box 217, 7500 AE Enschede, the Netherlands

## Abstract

*In this paper, the performance and characteristics of the execution of various join-trees on a parallel DBMS are studied. The results of this study, are a step into the direction of the design of a query optimization strategy that is fit for parallel execution of complex queries.*

*Among others, synchronization issues are identified to limit the performance gain from parallelism. A new hash-join algorithm, called Pipelining hash-join is introduced that has fewer synchronization constraints than the known hash-join algorithms. Also, the behavior of individual join operations in a join-tree is studied in a simulation experiment. The results show that the Pipelining hash-join algorithm yields a better performance for multi-join queries. Also, the format of the optimal join-tree appears to depend on the size of the operands of the join: A multi-join between small operands performs best with a bushy schedule; larger operands are better off with a linear schedule. The results from the simulation study are confirmed with an analytic model for dataflow query execution.*

## 1 Introduction

During the last years much attention has been paid to the development of parallel DBMSs. Using special-purpose hardware has shown not to be successful; instead, a parallel DBMS running on general-purpose shared-nothing hardware appears to be the right choice [DeG90]. Also, various query processing strategies have been implemented: dataflow query processing appears to be superior to control-flow scheduling [DGS90, TeB91]. Therefore, this paper studies query processing in a general-purpose, shared-nothing, dataflow architecture.

Teradata [Ter83], GAMMA [DGS90], Bubba [BAC90], HC16-186 [BrG89], and PRISMA [Ame91] are examples of parallel DBMSs that actually were implemented. Each of these systems exploits some sort of parallelism to speed up query execution. Within a query, inter-operator and intra-operator parallelism can be discriminated [ScD90, WAF91, WGA89]. Orthogonal to this distinction, pipelining can be contrasted to (pure) horizontal parallelism. This last type is called parallelism here, like in many other papers. *Intra-operator* parallelism is the primary source of parallelism in the projects mentioned above. This type of parallelism is well-understood now, and using it, efficient execution strategies can be found for simple queries The Wisconsin Benchmark [BDT83], which consists of simple queries on large volumes of data is used to describe the performance of a system [BrG89, DGS90].

A dataflow architecture, however, offers the possibility to also exploit *inter-operator* parallelism and pipelining by allocating different relational operations to differ-

ent (sets of) processors. The potential of using different types of parallelism for one query, turns query optimization into a difficult problem, that cannot be solved using conventional query-optimization techniques. So far, little has been done in this research area although it is identified to be important for the further development of parallel DBMSs [BAC90, DeG90]. The query optimizers for most parallel DBMSs are based on the theory developed in [SAC79], however, this theory is not particularly fit for parallel dataflow query processing. For example, only *linear* query trees are considered, although this class of trees does not necessarily include the optimal one for a parallel environment.

In a first attempt to understand the effect of various query tree formats, [ScD90] studies the behavior of right-deep and left-deep linear query trees for multi-join queries. It is concluded in that paper, that right-deep scheduling has performance advantages in the context of GAMMA. In [Gra90] it is shown how arbitrarily shaped query trees can be parallelized using the "exchange" operator, which splits a (part of) a query tree into a number of subtrees that can be executed in parallel. Although that paper makes clear that certain query trees can be parallelized, it does not solve the problem of which (type of) query tree performs best.

In this paper, we study the execution of multi-operation queries. The ultimate goal of this study is the design of a query optimizer for a parallel DBMS. Relational multi-join queries are used as an example, because the join is an important, and expensive relational operation. An outline of the path, we want to follow for this research is as follows: Query optimization comes down to selecting an execution strategy with low costs [Kui91]. Because searching the entire space of possible strategies is not feasible, most query optimizers are heuristic [BoR88, JaK84]. Heuristics are based on *insight* in the essentials of query execution. So, to design a heuristic query optimizer for a parallel DBMS, it is essential to *understand* the behavior of execution strategies for a query on a parallel DBMS. Modeling is a way to gain insight in the essentials of parallel query execution. Two approaches to modeling were used in our study: simulation and analytical modeling. The resulting knowledge, should eventually lead to the formulation of query optimization heuristics. It should be emphasized that we do not aim at a detailed quantitative model for a parallel DBMS, but rather at a simple, understandable framework that, by its nature, yields *insight* in the modeled phenomenon. However, neither intuition, nor the sort of model that is presented here can *validate* the heuristics that they yield, and so these heuristics have to be validated against a real dataflow DBMS, or a detailed simulation if

the former is not available. The methodology described in this paragraph, is similar to the methodology that is common in science: Scientists try to understand natural phenomena by modeling them. Subsequently, hypotheses are formulated, that are based on the model, and these hypotheses are validated against reality by experimentation.

This paper describes a first step on the path outlined above. First, the results of a simulation study are described. This work resulted in the proposal of a new join-algorithm that is fit for dataflow query execution. Also, the execution characteristics of multi-join queries were studied. An attempt to fully understand the results of this simulation study led to the derivation of an analytical model. The first results of this analytical model confirm the results of our simulation study. In the near future, we want to follow the path by extending the analytical model. That step should lead to the formulation of query optimization heuristics.

The research reported in this paper is carried out in the context of PRISMA/DB [Ame91,KAH87,WGA89]. PRISMA/DB is a parallel, Main-Memory, relational DBMS that runs on a 100-node shared-nothing architecture. The fact that PRISMA is a *Main-Memory* system plays an important role in our research. The price of primary memory has fallen sharply during the last years. As this trend is expected to continue, an interesting question arises: How can huge amounts of memory be used? In this study this question is specialized into: Can a very large primary memory yield performance gain in a DBMS? Therefore, we are willing to accept using large amounts of memory, if performance gain is expected in return.

The remainder of this paper is organized as follows: The next section describes dataflow query execution in a main-memory environment. Section 3 presents simulation results. Section 4 introduces the analytical model and its elaboration for join operations and join trees. The last section summarizes and concludes the paper.

## 2 Dataflow Query Execution

A Main-Memory parallel DBMS running on shared-nothing hardware has the following features: The hardware consists of a number of processors that can communicate via a message-passing network. Each processor hosts part of the base-data. A processor can access its part of the base-data directly. If a processor wants to access the data stored on another processor, it has to be sent to it via the network.

A query on a relational database can be represented as a dataflow graph. The nodes of such a graph represent eXtended Relational Algebra operations [WiG90]. Each processor can run one or more operations processes. In this paper, we want to study inter-operator parallelism, and therefore each operation process is assumed to have a private processor. Operation processes evaluate XRA-operations on their local data, or on tuple streams that are sent to them via the message-passing network. The result of the evaluation of an XRA-operation consists of a (multi)set of tuples. Such a result can either be stored locally, in which case it can be accessed by the local processor later on; or it can be sent to one or more other operation processes. In the last case the sending and the receiving operation processes can run concurrently, forming a pipeline.

Network transport of tuples is modeled as follows: To transport a tuple from a process to another, remote process, first, it has to be “wrapped” and put on the network hardware by the sending operating system, then, it is sent over the network, and finally, it has to be retrieved from the network and “unwrapped” by the receiving operating system. So, sending a tuple over the network implies CPU costs on the sending and receiving processor, and actual transmission, which implies a delay. In general, the CPU costs involved, appear to be the limiting factor, and, therefore, the *rate* in which tuples are transported over the network is determined by the capacity of the CPUs that send and receive the tuples and not by the capacity of the network hardware. So, tuple transport is modeled in terms of CPU costs on two processors [DeG90] and a constant transmission delay.

## 3 Simulation of dataflow query execution

This section describes the results of a simulation study. First, the simulation program used is described, then we study join algorithms, and finally we describe a simulation of multi-join queries.

### 3.1 The simulator

To study the execution characteristics of a query, a simulator for parallel query execution was developed. The input to the simulator is a schedule for a query. The output consists of a diagram for each operation process that was used in the schedule. These diagrams plot the processor utilization (on behalf of that operation process) against the time. The (horizontal) time-axis can be scaled. Figure 1 shows an example of such a diagram.



Figure 1: Sample output of the simulator

The simulator models local processing and network transport of tuples. The local processing model uses simple cost formulas for relational operations. Network transport of tuples is modeled as CPU-costs on the sending and receiving processor according to the description above. The simulator is parameterized with the costs of simple operations on tuples. The parameter values that were used in this paper are measured from the PRISMA/DB.

### 3.2 Join Algorithms

The choice of a join algorithm influences the execution characteristics of a join query in different ways.

Firstly, the processing, I/O, and communication costs are influenced. Schneider and DeWitt [ScD89] give an overview of well-known join algorithms and evaluate their performance for simple join-queries by experimentation. Hash-join algorithms are shown to be the most efficient ones for equi-joins. Therefore, in our paper, only Hash-join algorithms are considered.

Secondly, the synchronization between the joins that participate in a more complex join query, is determined by the join algorithm used. In this section, the synchronization requirements of a well-known Hash-join algorithm are

studied. Because those requirements are too tight to allow considerable performance gain from pipelining, a new main-memory hash-join algorithm is proposed that has fewer synchronization requirements [WiA90].

The known Hash-join algorithms, Grace hash-join, Simple hash-join, and Hybrid hash-join are disk based, and they only differ in the way disks are used. Therefore, only one Main-Memory version of these algorithms is dealt with in this paper. This algorithm is called Simple hash-join here.




Figure 2: Simple hash-join and Pipelining hash-join algorithm

### Simple hash-join

The Simple hash-join algorithm consists of two phases (See Figure 2). In the first phase, one entire operand is read into an in-memory hash-table. In the second phase, the tuples of the other operand are read one by one, each tuple is hashed and compared with the tuples in the corresponding bucket in the hash-table of the first operand. If a match is found, an output tuple is produced. This algorithm is asymmetric in its operands, although the join-operation is conceptually symmetric. The result is only formed during the second phase of the algorithm.

### Pipelining hash-join

The Pipelining hash-join algorithm (See Figure 2), aims at producing its output tuples as early as possible in the process of calculating the join, without decreasing the performance of the join operation itself. During the join process a hash table for both operands is built. The join process consists of only one phase. As a tuple comes in, it is first hashed and used to probe that part of the hash table of the other operand that has already been constructed. If a match is found, a result tuples is formed. Finally, the tuple is inserted in the hash table of its own operand. When the last tuple of one of the operands is processed, the join process can stop building a hash-table for the other operand, because this hash-table will not be used any more. Keeping this last feature in mind, it is easy to see that the Pipelining hash-join degenerates to a Simple hash-join when one operand is available to the join process entirely, before the first tuple of the other operand arrives. The Pipelining Hash-join algorithm is symmetric in its operands.

### Evaluation of the Simple hash-join and the Pipelining hash-join

Figure 3 shows the execution characteristics of the Simple hash-join processes, and of the Pipelining hash-join processes in a four-way multi-join  $((\sigma A \bowtie \sigma B) \bowtie (\sigma C \bowtie \sigma D))$  as visualized by the simulator. The figure shows the join tree for this query; two join symbols in each tree are replaced by the diagrams of the corresponding join processes. The time axis in all diagrams is scaled in the same way.




Figure 3: Simple Hash-join and Pipelining Hash-join

The join processes read their *input* from selection processes, that produce output at a limiting rate. The Pipelining Hash-Join makes a faster start than the Simple Hash-Join, because tuples belonging to both operands can be processed right from the beginning. Also, the Pipelining Hash-Join starts producing *output* earlier than the Simple Hash-Join. So, the consumer of the result of the Pipelining Hash-Join can start earlier than the consumer of the Simple Hash-Join. As a result, the response time of the evaluation with the Pipelining hash-join is better.

The CPU-utilization of the Pipelining Hash-Join is increasing in the time. This is caused by the increasing probability to find matching tuples as the hash-tables are filled. For join-operations that are higher up in a join tree this effect is enlarged by the fact that the operand tuples arrive at the join process with increasing rate.

The difference in synchronization requirements described above shows that the Pipelining hash-join allows more inter-operator pipelining than the Simple hash-join, and so it fits more naturally in a dataflow execution model.

It is the asymmetry in the Simple Hash-join algorithm that explains the difference in performance between left-deep, and right-deep scheduling reported in [ScD90]. Using a symmetric algorithm, like the Pipelining Hash-join yields the same performance for any linear join-tree. In the next section, the behavior of linear join-trees, and other join-tree formats is studied.

### 3.3 Multi-join queries

In this Section, the trade-offs of using differently structured query trees for the execution of multi-join queries are discussed. Also, the performance of the Simple, and the Pipelining hash-join in multi-join queries are compared. Figure 4 shows a join-query with a linear and a bushy join tree for that query.




Figure 4: A bushy and a linear join tree

## Trade-offs in join tree formats

First, some terminology is introduced. The term *hop* is used for the transmission from one join operation to its parent (the operation consuming its output) or its operand. The *termination delay* over one or more hops is the difference in termination time of the adjacent join operations. The term *delay* is used as a short-hand for termination delay.

Various types of nodes in a join tree can be identified.

- The leaf-nodes in a tree have two base-relations as operands. These base-relations are available to the join process immediately.
- The intermediate nodes in a linear tree and some intermediate nodes in bushy trees have one base relation and one intermediate result as operands. The base-relation is available to the join process immediately, but the join process has to wait for the other operand to become available from the previous join-operation.
- Bushy join trees contain join-processes that have two intermediate results as operands. Such a join process has to wait for both operands, and, therefore, the join process does not start immediately.

Having to choose a join tree for a join query, we are faced with the following trade-off: The join processes in a linear join tree can all start immediately hashing their base-relation operand; in this way, they fill the time waiting for the other operand. On the other hand, a linear join tree contains the longest possible pipeline, causing a large number of delays on top of each other. The pipelines in a bushy join tree are shorter than the ones in a linear one, but some intermediate join processes have to wait for both their operands, what may lead to large delays. In the next section, an experiment is described that shows that the optimal format of the join tree for a multi-join depends on the size of the operands.

## An experiment

To study the execution characteristics of various join trees, a join between 16 relations that have equal numbers of tuples, matching one tuple in each operand to exactly one tuple in another operand, is studied. All possible join trees for this query yield the same amount of joining and data communication costs. Also, the individual joins in the query are equal. So, any differences in response time are caused by the sort of join tree that is used only.

In four sub-experiments, the 16-way join described above is evaluated for operands with resp 1000, 5000, 10000, and 50000 tuples. The response times of those queries are measured with the simulator for a linear, and a completely bushy tree. The response times are shown in Table 1.

		1000	5000	10000	50000
Pipelining hash-join	bushy	3.9	16.2	31.1	149.0
	linear	6.8	13.9	23.0	95.4
Simple hash-join	bushy	5.5	22.5	43.8	214.0
	linear	7.0	15.4	25.8	109.5

Table 1: response times

Firstly, we want to remark, that the pipelining algorithm outperforms the Simple hash-join in all cases. The difference in performance is larger for bushy scheduling than for linear scheduling. This is caused by the fact that the Pipelining hash-join degenerates to the Simple hash-join in join operations that are relatively close to the root of a linear tree. In the remainder of this section, only the schedules using the Pipelining hash-join are considered.

Apparently, the bushy scheduling performs better for small operands and linear scheduling is better for large operands. Figures 5 and 6 show the execution characteristics of linear and bushy query trees for small and large operands. These figures show the join trees that were used. Some join symbols are replaced by the diagrams of the corresponding join-processes. In each join tree, the time-axis is scaled to the response time of the corresponding query.




Figure 5: Execution characteristics of linear join trees

From the diagrams in Figure 5, we see that in a *linear* tree, there is a *constant* delay over the pipeline between two adjacent joins. This delay does not depend on the size of the operands. Some diagrams show two distinct phases in the processing of the join. The first phase is the construction of a hash-table for the base-relation operand. The other phase is joining the other operand to this hash-table. In the lower most join operation these two phases are mixed, as described by the Pipelining hash-join algorithm. The further up in the join tree, however, the longer the join-operation has to wait for its second operand. At a certain point, the complete hash-table for the base-relation operand is ready before the first tuple of the second operand arrives. From this point on, the Pipelining hash-join behaves similar to the Simple hash-join. The point at which the two operands of the join are processed completely separately is reached earlier for small operands than for larger ones.

The diagrams for the bushy tree lead to the following observations. The leaf-nodes show the same characteristics as the leaf-node of the linear tree. The delay over one hop is larger than in case of a linear tree, because neither operand is directly available. Within one query each hop yields approximately the same delay. Moreover, a closer look at the characteristics shows that the delay over one hop is *proportional* to the size of the operands for operands of 1000 tuples or more. This proportionality is a surprising

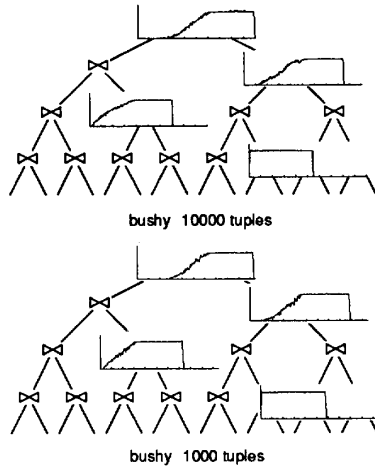


Figure 6: Execution characteristics of bushy join trees

result that cannot be accounted for intuitively.

The difference in termination delay between linear and bushy trees can explain the fact that bushy trees work better for small operands, and linear trees for larger operands. As the size of the operands grows, the termination delay, and so the response time of a bushy tree grows proportionally. For linear trees, on the other hand, the termination delay does not grow, and therefore its contribution to the total response time decreases with increasing operands. So, at a certain operand size, the linear trees outperform the bushy ones. With the parameter setting that was used in our simulation experiment, the break-even point lies at about 2000 tuples.

The next section introduces a mathematical model for dataflow query section that can explain the simulation results for bushy join-trees.

#### 4 An analytical model for dataflow query execution




Figure 7: Dataflow query execution

In this section, an analytical model for dataflow execution is developed. The general model can be specialized to model specific relational operations. This analytical model can explain the surprising proportional termination delay

in bushy query trees, that was found in the previous section.

#### 4.1 Some preliminaries Resources in the model

The model describes the *rates* at which tuples are transported and processed in a dataflow system. Also, the utilization of the processors participating in the dataflow system is modeled. Because, as described above, the bandwidth of the message-passing network is assumed to exceed the requirements of the application, the utilization of this hardware is not modeled. This paper only deals with retrieval, and in a Main-Memory context, retrieval does not need any disk-accesses. So, there is no need to model secondary storage either. The only resource that has to be taken into account now is the CPU. The resulting model is simple and consequently powerful: a complete analysis is possible for some classes of queries.

#### Modeling discrete phenomena

Tuples are discrete entities. Our model, however, is continuous. A continuous model for a discrete phenomenon is possible if large numbers of events are described [Oku80]. The transition from a discrete to a continuous model eliminates the need to use probability theory; If, in a discrete model, there is a probability 0.5 that a tuple is generated, the continuous model will generate half a tuple. This way of modeling has generally been accepted in physics and biology, and can be used here without problems.

#### Entities and dimensions

The rate at which tuples are transported and processed, and the utilization of processors are modeled. To do so, the costs of certain operations are expressed. Tuple transport is expressed in number of tuples per time-unit. The processor utilization is dimensionless, and has a maximum of 1. The costs of operations are expressed in time-units per tuple. Consider as an example, an operation that processes tuples at rate  $x$  tuples/time-unit. The processor spends  $A$  time-units for the processing of one tuple. The resulting processor utilization is  $Ax$  (dimensionless).

#### 4.2 Definition of a dataflow model

Figure 7 summarizes the essentials of a dataflow operation. The large box in this figure represents the processor; the small box represents the operation process. Data is sent to the operation process at the bottom of the box; the result is sent away at the top of the box.

#### Terminology

Each operation process has one or two operands. In this paper, only the join, which is a binary operation is considered, so there are always two operands. In Figure 7, each input stream contains two arrows: the first arrow indicates the rate at which tuples are *available* to the operations process, the second one represents the rate at which tuples are *processed* by the operation process. The arrow in the output stream indicates the rate at which result tuples are produced. The left column in Figure 7 shows the formalism used:

$a(t)$  is the rate at which tuples of a particular operand are available to an operation process at time  $t$ .

$x(t)$  is the rate at which tuples of a particular operand are processed by an operation process at time  $t$ .

$w(t)$  is the processor utilization at time  $t$ .

$r(t)$  is the rate at which tuples are produced at time  $t$ .

The functions  $a(t)$  and  $x(t)$  can be labeled with a subscript to indicate which operand is meant.

A query in execution, consists of a number of communicating operation processes. Time  $t = 0$ , is used to indicate the starting point of the *entire* query. Some operation processes in the query may be idle at time  $t = 0$ .  $T$  is used for the time at which an operation process is ready.

### 4.3 Some relationships

From the description of dataflow execution, the following relationships can be deduced:

An operation process cannot process tuples in an operand before they have arrived at its input stream. This can be modeled as:

$$\forall t \text{ in } [0, T] : \int_0^t a(\tau) d\tau \geq \int_0^t x(\tau) d\tau$$

The processor utilization is a function of the rate at which the operand tuples are processed:

$$w(t) = \mathcal{W}(x_1(t), x_2(t)) \quad (1)$$

The rate at which tuples are produced is also a function of the rate at which the operand tuples are processed:

$$r(t) = \mathcal{R}(x_1(t), x_2(t))$$

If a CPU works at full capacity, its utilization is 1. Therefore,  $w(t)$  can never be larger than 1.

Our model discriminates between the rate at which operand tuples are available, and the rate at which they are processed. This is done, because these two rates may differ, if the operation process cannot keep up with the rate at which tuples are sent to it. This observation leads to the definition of two different modes in which an operation process can work.

**input-limited mode** Tuples are sent to the operation process at such low rate, that the operation process can keep up with this rate. Now,  $w(t) \leq 1$ , and  $x_j(t) = a_j(t)$  for both operands.

**CPU-limited mode** Tuples are sent to the operation process at such high rate, that the receiving processor cannot keep up with this rate, so  $w(t) = 1$  and there is an operand for which  $x_j(t) < a_j(t)$ .

This discrimination leads to the central equation in this paper:

$$\begin{array}{ll} x_j(t) = a_j(t) & \text{if process input-limited} \\ x_j(t) \text{ meets } w(t) = 1 & \text{if process CPU-limited} \end{array} \quad (2)$$

This equation is used to evaluate the behavior of an operation process. An outline of such an evaluation is as follows: Equation (1) expresses the CPU-utilization as a function of the rates at which operand streams are processed. In the input-limited mode of an operation process  $x_j(t) = a_j(t)$ , so evaluation of  $\mathcal{W}((a_1(t), a_2(t)))$ , and comparing the result to 1 (the maximal CPU-utilization) can reveal whether an operation process is input-limited, or CPU-limited at time  $t$ . If an operation process is input-limited, the rate at which operand streams are processed is clear,  $a(t)$ . If, on the other hand, the process appears to be CPU-limited, then solving equation  $\mathcal{W}((x_1(t), x_2(t))) = 1$ ,

for  $x_j(t)$  shows at what rate each operand tuple-stream is processed.

Knowing the functions  $x_j(t)$  and mapping  $\mathcal{R}$ , the rate at which tuples are produced by an operation process,  $r(t)$ , can be calculated. The result of an operation process can be sent as input to another operation process. Those tuples are assumed to arrive at the receiving operation process with some delay<sup>1</sup> at the rate at which they are produced by the producing operation process. So, then function  $a(t)$  for a consumer is known, and we are in the position to evaluate the behavior of this consumer process.

Summarizing, the model maps the rate at which operand tuples are available, to the rate at which result tuples are produced. To describe a query tree, the result of the evaluation of one operations can be used as input to a next one.

In the remainder of this section, the model developed above is specialized to describe the Pipelining hash-join. After that, the results of this evaluation are used to study bushy join-trees. In all cases, the goal is full characterization of the participating operations in terms of  $x(t)$ ,  $w(t)$ ,  $r(t)$ , and  $T$ , given the rate at which operand tuples are available,  $a(t)$ .

### 4.4 Pipelining Hash-join

In this section, it is assumed that the operand tuples from both operands are available at a non-limiting rate, so, the operation process only processes in its CPU-limited mode. The more general case, in which the join-process works both input-limited, and CPU-limited, is dealt with afterwards. Both operands are equal in size: they contain  $n$  tuples. The selectivity of the join operation is assumed to be  $g$ : the result contains  $gn^2$  tuples.

From the description of the Pipelining join algorithm, it is clear that the join algorithm processes tuples from both operands at the same rate,  $x(t)$ . The goal of this section is finding  $x(t)$ , and deriving  $T$ , and  $r(t)$  from  $x(t)$ .

During the join operation, each operand tuple has to be made available to the join process, its hash-value has to be calculated, it has to be inserted in a hash-table and it has to be compared to the tuples in the corresponding bucket of the other operands hash-table. These costs are assumed to be constant during the join-process ( $A$ )<sup>2</sup>. If a match is found, a result tuple is generated. The costs associated with producing a tuple are assumed to be constant ( $S$ ). The distinction between work dedicated to processing operand tuples and to generating result tuples is essential. Through this distinction, insight is gained in how an operation process maps its input streams to an output stream.

#### The model

The tuples in both operands are processed at the same rate,  $x(t)$ . So, the amount of work the processor spends on processing input tuples is equal to:

$$2Ax(t)$$

<sup>1</sup>The transmission delays can easily be handled in our model, and their influence on the results is simple. The formalism is complicated by using them however, and also, the transmission delay is assumed to be small compared to the time needed to evaluate a relational operation. Therefore, we choose not to incorporate them in the model in this paper.

<sup>2</sup>Actually these costs are increasing slightly during the join process, due to the fact that hash-buckets are filled. Using a good hash-table though, minimizes this increase.

There is a factor 2 in this expression, because tuples from both operands are processed.

The amount of work spent on generating the result is calculated as follows: The probability that a tuple, arriving at time  $t$ , produces a match is proportional to the number of tuples that have already arrived in the other operand. As the number of tuples that have arrived at the operation process at a certain time is equal to the integral over the rate at which the tuples arrive, the number of result-tuples that is formed upon the arrival of one tuple at time  $t$  is equal to:

$$\rho \int_0^t x(\tau) d\tau$$

Using this expression, the amount of work spent on generating the result can be formulated as:

$$2\rho Sx(t) \int_0^t x(\tau) d\tau$$

The CPU-utilization is equal to the sum of the amount of work spent on processing the input, and the amount of work spent on generating the output.

$$w(t) = 2Ax(t) + 2\rho Sx(t) \int_0^t x(\tau) d\tau \quad (3)$$

The join operation processes in its CPU-limited mode, so equation (2) can be specialized to:

$$x(t) \text{ meets } 2Ax(t) + 2\rho Sx(t) \int_0^t x(\tau) d\tau = 1 \quad (4)$$

### Finding the rate at which operand tuples are processed

We are now ready to find  $x(t)$  from equation (4). This integral equation can be solved using elementary calculus [Wil91]:

$$x(t) = \frac{1}{2\sqrt{\rho St + A^2}} \quad (5)$$

The bottom row of Figure 8 shows diagrams that plot  $x(t)$  against  $t$ . As expected,  $x(t)$  decreases in the time, because, as the join processes advances, more effort has to be spent on generating result tuples, because the hash-tables are filled.

### Termination of the join process

The join operation is ready at time  $T$ . At this time,  $n$  tuples per operand have been processed:

$$\int_0^T x(\tau) d\tau = n$$

Substitution of Equation 5 and solving this equation for  $T$  yields:

$$T = 2An + \rho Sn^2 \quad (6)$$

This result is reasonable:  $2n$  operand tuples and  $\rho n^2$  result tuples have to be processed. This processing costs  $2An + \rho Sn^2$  units of time. If the CPU-utilization is equal to 1 during the entire join process, the join process will end at  $2An + \rho Sn^2$ .

### The output stream

The rate at which result tuples are produced can be derived from Equation 4:

$$r(t) = 2\rho x(t) \int_0^t x(\tau) d\tau = \frac{1 - 2Ax(t)}{S}$$

Substitution of (5) yields:

$$r(t) = \frac{1}{S} \left\{ 1 - \frac{A}{\sqrt{\rho St + A^2}} \right\} \quad (7)$$

The top row of Figure 8 shows diagrams that plot  $r(t)$  against the time.  $r(t)$  is increasing in time, because the probability of finding a match in the join-process increases with the time.

### Two join-processes with different selectivity

Figure 8 shows some diagrams of the characterization of




Figure 8: Pipelining joins with different selectivities

two different join-processes<sup>3</sup>. Two join-operations are illustrated: one in which the selectivity of the join was chosen to be  $1/n$  so that the result contains  $n$  tuples, and another in which the selectivity of the join is  $2/n$ , so that the result contains  $2n$  tuples. The first join-operation is in the first column of diagram in the figure, and the second join-operation is in the second.

The topmost diagrams show the rate at which result tuples are generated as a function of the time. This rate is increasing in the time due to increasing probability of finding a match. At the very beginning of the join-operation  $r(t)$  is equal to zero, due to the fact that no tuples to form a match with have arrived yet. As expected, the second join-operation produces tuples at a higher rate than the first one.

The middle diagrams show the processor-utilization as a function of the time. As the join-operations are CPU-limited during the entire operation, the processor utilization is equal to 1. In these diagrams, an additional curve

<sup>3</sup>Part of the symbolic manipulation and the generation of plots of the results of this manipulation was carried out using the symbolic manipulator Maple [CGG88].

shows what portion of the CPU-effort is spent on processing input tuples (area below the curve), and what portion is spent on generating output (other area). We see that the less selective join-operation spends a larger portion of its effort on generating output, and that in both cases the amount of work related to generating output increases, at the expense of processing input.

The bottom diagrams show the rate at which input tuples are processed. As expected this rate is decreasing, and this effect is stronger for second join-operation.

The model developed in this section, yields an analytical expression for the rate in which output tuples are produced (Equation 7). In the next section the output stream of a single join-operation is used as input to a next join-operation: in this way join trees are studied.

#### 4.5 Symmetric Bushy Join trees

The behavior of symmetric bushy join trees is analyzed. Figure 4a shows a symmetric tree for an 8-way join. It is assumed that the operands are equal in size (each operand contains  $n$  tuples) and that the join operations match one tuple in their left operand to exactly one tuple in their right operand, so the selectivity ( $q$ ) of the join-operations is equal to  $1/n$ . Each join-operation has a private processor. The bushy join tree in the simulation experiment that was described in Section 3 corresponds to the sort of join trees that are modeled here. It is clear, that the join operations that have two base relations as operands, all have the same execution characteristics. These join operations are called level0 joins. The join operations that join the results of level0 joins again have the same characteristics. They are called level1 joins. In the same way, level2, level3, and even higher levels can be defined.

##### The model

It is assumed, that the base operands are available to the level0 joins at non-limiting rates. Therefore, the characteristics of these joins are as described in Section 4.4. To describe the other levels, some notation conventions are needed. In this section, subscripts are used to indicate the level of the join-operation. Due to the symmetry of the problem, we do not need to discriminate between the two input streams of one join-operation. So,  $x_i(t)$  denotes the rate at which tuples of one level1 join-operand are processed.

The level0 joins operate in CPU-limited mode from the beginning on. Higher level joins, however, are expected to show an increase in their CPU-utilization: they start input-limited, and after some time they switch to CPU-limited. The following symbols are used to describe this:

$\theta_i$  The time at which a CPU executing a level $i$  join is saturated. So, at this time, the join process switches from its input-limited to its CPU-limited mode.  $\theta_0$  is equal to 0.

$\mathcal{H}_i$  The number of tuples that have been collected in the hash-table of an operand of a level $i$  join at time  $\theta_i$ .

$\mathcal{H}_0$  is equal to 0 and  $\mathcal{H}_i = \int_0^{\theta_i} x(\tau) d\tau$ .

Result tuples from one level are sent as input to a join-operation in the next level. So,

$$a_{i+1}(t) = r_i(t) \quad (8)$$

Now, we can derive the model for a join-operation at level $i$  of a bushy join tree. Equation (8) is used to characterize the input rate, and the central equation of this paper (2) is used to model the operation process. Furthermore, an expression for CPU-utilization (3) for the Pipelining join is derived in the previous section. The combination of these three equations, and the definition of  $\theta_i$ , yields the model for level  $i$  of a bushy join tree:

$$\begin{aligned} x_i(t) &= r_{i-1}(t) \text{ if } 0 \leq t < \theta_i \\ x_i(t) \text{ meets } & \begin{cases} 2Ax_i(t) + 2\frac{S}{n}x_i(t) \int_0^t x_i(\tau) d\tau = 1 \\ \text{if } t \geq \theta_i \end{cases} \end{aligned}$$

These equations can be solved explicitly for  $x_i(t)$  [Wil91]:

$$\begin{aligned} x_i(t) &= r_{i-1}(t) \text{ if } 0 \leq t \leq \theta_i \\ x_i(t) &= \frac{1}{2\sqrt{\frac{S}{n}(t-i\delta) + A^2}} \text{ if } \theta_i \leq t \leq T_i \\ r_i(t) &= \frac{1 - 2Ax_i(t)}{S} \\ T_0 &= n(2A + S) \\ T_i &= T_0 + i\delta \\ \delta &\text{ is proportional to } n \\ \theta_i &= \theta_1 + (i-1)\delta \quad i \geq 1 \\ \theta_1 &\text{ is proportional to } n \\ \mathcal{H}_i &= \mathcal{H}_j \quad i, j \geq 1 \end{aligned}$$

This solution looks rather complex, but it can be interpreted in the following way.

- Two phases can be discriminated in a join-process: The *startup*-phase, in which the join-process does not saturate its processor, and the *main*-phase, in which the join process is CPU-limited. In a level0 join the startup-phase has length 0 (so, there is no startup phase).
- The main phases of subsequent join levels are similar, apart from a translation  $\delta$  in time. This implies that the main phase of each subsequent level starts and ends  $\delta$  time-units after its predecessor. The important thing to note here, is the fact that each subsequent level terminates  $\delta$  time-units later. As  $\delta$  is proportional to the number of tuples in an operand, this fact confirms our former observation that the termination delay in a bushy join-tree is proportional to the number of tuples in the operands.
- The startup phase of a join-operation takes longer for higher levels, but the number of tuples that are processed ( $\mathcal{H}$ ), and consequently the amount of work that is done during startup is equal for each level.

##### Examples

Figure 9 shows diagrams that plot  $w(t)$  against the time for level0 through level2. Similar to the diagrams for  $w(t)$  in Figure 8, an additional plot in the diagrams shows what portion of the CPU-time is spent on processing the input. Figure 9 has two columns of diagrams: the first column is of a 16-way join between operands of 1000 tuples; the second one shows the same join with operands of 1500 tuples.



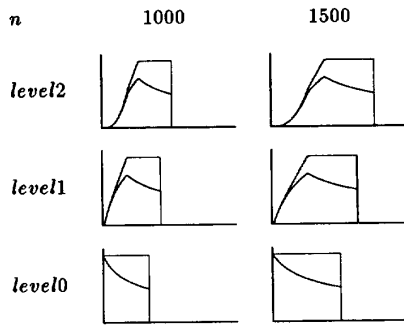


Figure 9: Join processing in a bushy tree

The time axis of all diagrams is scaled in the same way. Comparison of the diagrams in one column shows that the delay between subsequent levels is constant. Comparison of all diagrams shows that the termination delay is proportional to the number of tuples in one operand.

### Optimization and scheduling of multi-join queries

Although the model for join trees has to be extended to cover general join trees, it can now already be indicated how the sort of results that the model yields (with the results for symmetric bushy trees as an example), can be used for query optimization, and query scheduling. Apparently, the results not only indicate the *amount* of work that has to be done on behalf of one relational operation; they also give an indication of *when* this work has to be done and how busy a processor will be. The following examples illustrate how this scheduling information can be used.

- A query optimizer, having to select an execution strategy for a query, can use both the amount of work that has to be done and the timing information. It is possible that a more expensive schedule (in terms of total processing costs) has very good timing characteristics, so that its response time is very good. If response time is the important figure in the system, such a schedule should be selected.
- The knowledge about when processors are busy can be used by a scheduler: A processor that is assigned to a relational operation can be used for other purposes (possibly another relational operation) during the time that it is idle.
- The model for bushy query trees shows that in the execution of a join, a startup, and a main phase can be discriminated. If the scheduling of a join at level  $i$  is postponed until time  $i\delta$ , the main phase of the join is left the same, and the startup phase uses the processor at full capacity during a shorter period of time. Therefore, postponing the scheduling of a level  $i$  join until time  $i\delta$  does not affect the response time of the entire query. In Figure 10, the characteristics of a level2 join are shown. The left diagram shows the characteristics of a join that is scheduled immediately after query startup, the right one shows the characteristics when the join is scheduled at time  $2\delta$ .




Figure 10: Scheduling of join processes

Although a limited type of join trees was modeled in this section, we feel that the model increased our insight in the working of the pipelining join algorithm, and its cooperation with its producers and consumers. Specifically, the fact that the model can predict when the higher level join operation need CPU capacity is encouraging. Also, the fact that postponing the scheduling of certain join-operations to some extent does not influence the response time to the entire query can be used easily in a scheduler: two different relation operations can be scheduled subsequently to one processor.

## 5 Conclusions and future work

The work reported in this paper is a part of our research on query optimization strategies for a parallel dataflow DBMS. In the introduction to this paper, a methodology to do this research was outlined. It was illustrated that gaining understanding of parallel dataflow query processing is an essential step. This understanding should be used to design a heuristic query optimizer for a parallel dataflow DBMS.

The work reported in this paper was introduced as a step into the direction of understanding parallel dataflow query execution. Looking back, we are faced with some questions: “What insight was gained from our study?”, “How can this knowledge be used?”, and “How can this knowledge be validated?”. These questions are now answered in turn.

### What did we learn from our study?

- The simulation study showed that different aspects of the algorithms that are used for relational operations in a query tree are important. Apart from, of course, the CPU-costs of an algorithm, also its synchronization with the processes that produce and consume its input and output is important to yield a good performance. It was shown that the well-known Simple Hash-join algorithm has synchronization requirements that are too tight to allow performance gain from pipelining. A new hash-join algorithm, the Pipelining Hash-join, was proposed that is expected to give good performance in a dataflow system. Algorithms for other relational operations can be studied in the same way. Pipelining algorithms are possible for many relational operations.
- The simulation study of various join-tree formats gave insight in the behavior of individual join processes in relationship to their position in a join tree. It was shown that the time at which a join-operation can start processing depends on the position of the join operation in the join tree, and on the sizes of the operands. Regular linear, and bushy trees were studied extensively. Other join-tree formats need additional study.
- The mathematical model which was developed in this paper confirms some of the simulation results. Also, the model can predict the effect of changes in the

scheduling of join-processes, as indicated in the previous section

*How can this knowledge be used?*

The gained insight can be used in several ways. Firstly, pipelining algorithms should be used in dataflow systems. Secondly, a query optimization design, based on the ideas developed in this paper becomes feasible. Knowing the CPU-costs, and the delays that are incurred in a join-tree, the response time to a join-query can be calculated. This cost evaluation can be used in combination with known query optimization techniques, that search (part of) the space of possible execution strategies to find the cheapest one. Finally, the insight in the timing requirements can be used in a scheduler, as illustrated in Figure 10.

*How do we validate this knowledge.*

Currently, we are planning experiments on the latest version of PRISMA/DB.

Following the path of our research requires the study of more general join trees. Some preliminary work in this direction has been done and the results are encouraging. We plan to incorporate other relational operations in our model, and we want to study the effect of distributing individual relational operations. Although the model was only evaluated for a limited class of queries, we can now already make statements about the scheduling of operations in query trees. It should be emphasized that, as explained in the introduction, these statements have to be validated.

This paper is about parallel query execution. The concrete results of our study are worthwhile, and they probably eventually will be used to design new query execution strategies. We feel however, that apart from the concrete results, the approach to obtaining them also is a contribution to the database research. The experimental approach that is adopted from science, combined with mathematical modeling of the observed phenomena, is, to our opinion, a viable methodology to tackle certain problems in computer science.

## References

- [Ame91] P. America, ed., *Proc. PRISMA Workshop on Parallel Database Systems*, Springer-Verlag, New York-Heidelberg-Berlin, 1991.
- [BDT83] D. Bitton, D. J. DeWitt & C. Turbyfill, "Benchmarking Database Systems - A Systematic Approach," in *Proc. 9th VLDB Conference, Florence, Italy, October 31-November 2, 1983*.
- [BoR88] P. Bodorik & J. S. Riordon, "Heuristic Algorithms for Distributed Query Processing," in *Proc. Int. Symp. on Databases in Parallel and Distributed Systems, Austin, Texas, USA, December 1988*, 107 - 117.
- [BAC90] H. Boral, W. Alexander, L. Clay, G. Copeland, S. Danforth, M. Franklin, B. Hart, M. Smith & P. Valduriez, "Prototyping Bubba, A Highly Parallel Database System," *IEEE Transactions on Knowledge and Data Engineering* 2 (1990), 4-24.
- [BrG89] K. Bratbergsengen & T. Gjelsvik, "The Development of the CROSS8 and HC16-186 (Database) Computers," in *Proc. 6th IWDM, Deauville, France, June 1989*, 359 - 372.
- [CGG88] B. W. Char, K. O. Geddes, G. H. Gonnet, M. B. Monager & S. M. Watt, *Maple Reference Manual*, WATCOM Publications Limited, Waterloo, Canada, 1988.
- [DGS90] D. J. DeWitt, S. Ghandeharizadeh, D. A. Schneider, A. Bricker, H. Hsiao & R. Rasmussen, "The GAMMA Database Machine Project," *IEEE Transactions on Knowledge and Data Engineering* 2 (March 1990), 44-62.
- [DeG90] D. J. DeWitt & J. Gray, "Parallel Database Systems: The Future of Database Processing or a Passing Fad?," *SIGMOD RECORD* 19 (1990), 104 - 112.
- [Gra90] G. Graefe, "Encapsulation of parallelism in the Volcano Query Processing System," in *Proc. ACM-SIGMOD 1990, Atlantic City, NJ, May 23-25, 1990*, 102-111.
- [JaK84] M. Jarke & J. Koch, "Query Optimization in Database Systems," *Computing Surveys* 16 (June 1984), 111 - 152.
- [KAH87] M. L. Kersten, P. M. G. Apers, M. A. W. Houtsma, H. J. A. van Kuijk & R. L. W. vande Weg, "PRISMA: A Distributed Main Memory Database Machine," in *Proc. 5th IWDM, Karuizawa, Japan, October 1987*.
- [Kui91] E. van Kuijk, *Semantic Query Optimization in Distributed Database Systems*, PhD-Thesis, University of Twente, 1991.
- [Oku80] A. Okubo, *Diffusion and Ecological Problems: Mathematical Models*, Springer-Verlag, New York-Heidelberg-Berlin, 1980.
- [ScD90] D. A. Schneider & D. J. DeWitt, "Tradeoffs in Processing Complex Join Queries via Hashing in Multiprocessor Database Machines," in *Proc. 16th VLDB Conference, Brisbane, Australia, August 13-16, 1990*, 469-480.
- [ScD89] D. A. Schneider & D. J. DeWitt, "A Performance Evaluation of Four Parallel Join Algorithms in a Shared-Nothing Multiprocessor Environment," in *Proc. ACM-SIGMOD 1989, Portland, OR, May 31-June 2, 1989*.
- [SAC79] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie & T. G. Price, "Access Path Selection in a Relational Database Management System," in *Proc. ACM-SIGMOD 1979, Boston, USA*, 82 - 93.
- [TeB91] W. B. Teeuw & H. M. Blanken, "Control versus data flow in distributed database machines," Memorandum INF91-02, Universiteit Twente, Enschede, The Netherlands, 1991.
- [Ter83] Teradata Corporation, "Teradata, "DBC/1012 Database Computer Concepts and Facilities," C02-0001-00, 1983.
- [Wil91] A. N. Wilschut, "A model for the pipelining hash-join algorithm," Memorandum INF91-34, Universiteit Twente, Enschede, The Netherlands, 1991.
- [WiA90] A. N. Wilschut & P. M. G. Apers, "Pipelining in Query Execution," in *Proc. Int. Conference on Databases, Parallel Architectures and their Applications, Miami, USA, March 1990*.
- [WAF91] A. N. Wilschut, P. M. G. Apers & J. Flokstra, "Parallel Query Execution in PRISMA/DB," in *Proc. PRISMA Workshop on Parallel Database Systems, Noordwijk, The Netherlands, September 1990*, P. America, ed., Springer-Verlag, New York-Heidelberg-Berlin, 1991.
- [WiG90] A. N. Wilschut & P. W. P. J. Grefen, "PRISMA/DB1 XRA Definition," PRISMA Document P465, Philips Research Laboratories, Eindhoven, The Netherlands, 1990.
- [WGA89] A. N. Wilschut, P. W. P. J. Grefen, P. M. G. Apers & M. L. Kersten, "Implementing PRISMA/DB in an OOPL," in *Proc. 6th IWDM, Deauville, France, June 1989*, 359 - 372.