

**A Trace-Driven Comparison of Algorithms for
Parallel Prefetching and Caching**

Tracy Kimbrel, Andrew Tomkins, R. Hugo Patterson,
Brian Bershad, Pei Cao, Edward W. Felten,
Garth A. Gibson, Anna R. Karlin, and Kai Li

Technical Report 96-09-01

Department of Computer Science and Engineering
University of Washington

Department of Computer Science and Engineering, Box 352350
University of Washington, Seattle, WA 98195-2350 USA

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Parallel prefetching and caching	2
1.3	Comparing approaches	2
1.4	Summary of results	3
1.5	Related work	3
1.6	Organization of the paper	4
2	The algorithms	4
2.1	Theoretical model	4
2.2	Optimal prefetching rules	4
2.3	The <i>fixed horizon</i> algorithm	4
2.4	The <i>aggressive</i> algorithm	5
2.5	The <i>reverse aggressive</i> algorithm	5
2.6	Practical considerations	5
2.7	Implementations of the algorithms	6
3	Simulation framework	6
3.1	File access traces	7
3.2	Data placement and disk head scheduling	8
4	Results	8
4.1	Comparison with demand fetching	8
4.2	Fundamental differences	8
4.3	Application traces	9
4.4	Varying parameters	11
5	A new approach	13
5.1	Performance of <i>forestall</i>	14
6	Conclusions	15
A	Performance data: baseline measurements	17
B	Performance data: FCFS	23
C	Performance data: double-speed CPU	28
D	Performance data: varying cache size	29
E	Performance data: varying <i>aggressive's</i> batch size	33
F	Performance data: varying <i>reverse aggressive's</i> parameters	42
G	Performance data: varying <i>fixed horizon's</i> horizon	51
H	Performance data: <i>forestall</i> with a fixed fetch time estimate	56

A Trace-Driven Comparison of Algorithms for Parallel Prefetching and Caching

Tracy Kimbrel* Andrew Tomkins† R. Hugo Patterson‡ Brian Bershad* Pei Cao§
Edward W. Felten¶ Garth A. Gibson† Anna R. Karlin* Kai Li¶

Abstract

¹ High-performance I/O systems depend on prefetching and caching in order to deliver good performance to applications. These two techniques have generally been considered in isolation, even though there are significant interactions between them; a block prefetched too early reduces the effectiveness of the cache, while a block cached too long reduces the effectiveness of prefetching. In this paper we study the effects of several combined prefetching and caching strategies for systems with multiple disks. Using disk-accurate trace-driven simulation, we explore the performance characteristics of each of the algorithms in cases in which applications provide full advance knowledge of accesses using hints. Some of the strategies have been published with theoretical performance bounds, and some are components of systems that have been built. One is a new algorithm that combines the desirable characteristics of the others. We find that when performance is limited by I/O stalls, aggressive prefetching helps to alleviate the problem; that more conservative prefetching is appropriate when significant I/O stalls are not present; and that a single, simple strategy is capable of doing both.

1 Introduction

Recently there has been a great deal of interest in prefetching from parallel disks as a technique for improving the I/O performance of sequential applications. In this paper, we study prefetching and caching strategies for multiple disks in the presence of application-provided knowledge of future accesses. We compare the performance of four algorithms:

1. *Fixed horizon* is simple to implement, and has near-optimal performance when sufficient I/O parallelism is

available, but can be suboptimal in I/O-bound situations.

2. *Aggressive* is also simple to implement, is close to optimal for a single disk and for well-laid-out data on multiple disks, but can be suboptimal for multiple disks when the load on the disks is unbalanced.
3. *Reverse aggressive* is substantially more complex, but is provably close to optimal for all configurations in a uniform fetch-time model of disk accesses.
4. *Forestall* is a new algorithm, representing an attempt to combine the desirable characteristics of the other three algorithms.

Using trace-driven simulation on a collection of file access traces, we compare the performance of these algorithms assuming an environment in which a single process is running and full advance knowledge is available.

1.1 Motivation

Our work is motivated by recent advances in technology that have made magnetic disks both cheaper and smaller. As a result, parallel disk arrays have become an attractive means for achieving high performance from storage devices at low cost [15, 27, 23]. Independently accessible multiple disks offer the advantage of both increased bandwidth and reduced contention on individual disk arms. However, many applications do not benefit from this parallelism because their I/O accesses are serial. This problem is particularly severe for read-intensive applications. Write performance is less important as write behind strategies can mask update latency. Read-intensive applications that stall for I/O a significant fraction of their running time include text search, 3D scientific visualization, relational database queries, multimedia servers and object code linkers.

Many of these applications have *predictable access patterns* [24, 1, 17]. The ability to provide the file system with hints about future references has motivated research into the design of policies that use this information to reduce I/O overhead [24, 25, 7, 6]. The two key techniques that are enabled by detailed information about future accesses are deep prefetching and better-than-LRU cache replacement.

*Dept. of Computer Science and Engr., University of Washington, Seattle WA ({tracyk,bershad,karlin}@cs.washington.edu)

†School of Computer Science, Carnegie Mellon University, Pittsburgh PA ({andrewt,garth}@cs.cmu.edu)

‡Dept. of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh PA (rhp@cs.cmu.edu)

§Computer Sciences Dept., University of Wisconsin - Madison, Madison WI (cao@cs.wisc.edu)

¶Dept. of Computer Science, Princeton University, Princeton NJ ({felten,li}@cs.princeton.edu)

¹This paper originally appeared without the appendices in the *Proceedings of the Second ACM SIGOPS/USENIX Symposium on Operating System Design and Implementation*.

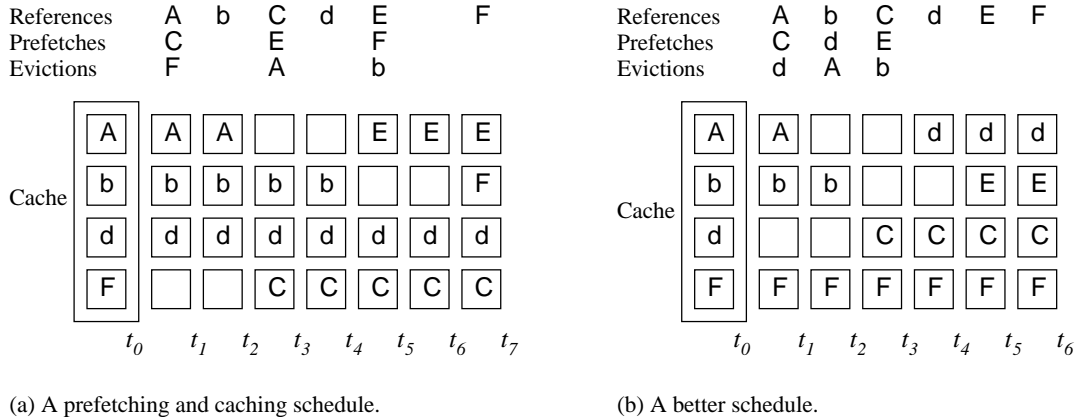


Figure 1: An example of prefetching and caching with advance knowledge of accesses for two disks. One disk holds blocks A, C, E, and F, and another disk holds blocks b and d. Cache size $k = 4$ disks and fetch time $F = 2$ computation steps.

This paper explores the tradeoff between aggressive prefetching and optimal cache replacement. The decision to prefetch requires that a buffer be reserved immediately, usually precipitating the replacement of cached data. This earlier replacement may result in an inferior replacement choice, which may actually increase the total number of fetches and degrade performance. Furthermore, in the multiple disk case, poor replacement choices can lead to load imbalance between the disks. The algorithms we study choose different points on the spectrum between aggressive prefetching (with possibly poorer cache replacements) and more conservative prefetching (with closer to optimal cache replacement).

The algorithms explored in this paper also differ in the degree to which they require advance knowledge. *Fixed horizon* exploits global knowledge the least while *reverse aggressive* exploits it the most. Using disk-accurate, trace-driven simulation, this paper’s results provide a measure of the potential benefit of using global knowledge.

1.2 Parallel prefetching and caching

Prefetching and caching are even more complicated in a system with multiple disks, not only because it is possible to do multiple prefetches in parallel, but also because appropriate cache replacement strategies can alleviate load imbalance among the disks. Since a disk can serve only one prefetch at a time, a set of blocks can be prefetched in parallel only if they reside on different disks.

In order to develop intuition for why cache replacement strategies can affect parallel prefetching performance, consider the example shown in figure 1. In this example, the cache holds four blocks and the application references one block per time unit. If the application wants to reference a block that is not present in the cache, the application must wait or *stall* until the block is present. Suppose that it takes two time units to fetch a block from disk, and that the fetches on each disk are serialized. Every fetch evicts some block from the cache; the evicted block becomes unavailable at the moment the fetch starts. The goal is to minimize the total time spent by the application.

The application references blocks in the sequence (A, b, C, d, E, F) , and the cache initially holds blocks A, b, d , and F . Blocks A, C, E and F are on one disk, and blocks b and d are on the other disk. The straightforward approach is to prefetch aggressively: always fetch the missing block that will be referenced soonest; evict the block whose next reference is furthest in the future; but do not fetch if the evicted block will be referenced before the fetched block. For small caches such as in this figure, the fixed horizon and aggressive algorithms both behave in this way.

Figure 1(a) shows the cache block changes using this approach. The total elapsed time for the sequence is 7 time units. Figure 1(b) shows another prefetching schedule on the same reference pattern that is faster by one time unit. On the first fetch, d is evicted rather than F , even though d is referenced earlier. This has the advantage of offloading one fetch from the heavily loaded disk (the one holding A, C, E , and F) to the otherwise idle disk (the one holding b and d). This change allows two fetches to proceed in parallel later, thus saving one time unit.

The example shows that the achievable I/O parallelism of multiple disks can be affected by cache replacement and data placement policies. These are the factors that are addressed by the reverse aggressive algorithm.

1.3 Comparing approaches

The *fixed horizon* algorithm is based on the second Informed Prefetching (TIP2) system of Patterson, Gibson *et al.* [25], which manages allocation of cache space and I/O bandwidth between multiple processes, only some of which are disclosing some or all of their future accesses. TIP2 is designed for the case in which sufficient I/O bandwidth exists to service the request stream without stalling on I/O. The *fixed horizon* algorithm, a restriction of TIP2 to a single hinting process, initiates a fetch for a missing block H references ahead of its reference, where H is the ratio between the average time it takes to read a block from disk and the minimum time it takes to consume a block of data. Patterson *et al.* showed that under the assumption of sufficient bandwidth,

this strategy eliminates stalls while placing little stress on system resources. However, *fixed horizon* will not look further than H references into the future for fetches to perform. This can cause it to stall on I/O when there is insufficient disk bandwidth.

In contrast, the *aggressive* and *reverse aggressive* algorithms are designed to take maximum advantage of any amount of I/O parallelism. They use knowledge of future accesses to minimize application elapsed time for both small and large numbers of disks. *Aggressive* prefetches as early as possible, provided that the prefetched block is needed by the application sooner than the block that it will replace. When insufficient bandwidth is available, in particular, it becomes more important to schedule prefetch requests to ensure that no bandwidth is wasted.

Reverse aggressive goes beyond *aggressive*'s use of future knowledge by attempting to balance disk workload through carefully selected replacement decisions. Previously, it was shown theoretically that on any access pattern known in advance, *reverse aggressive*'s elapsed time is close to optimal [16]. It is the only one of these algorithms with this theoretical performance guarantee. However, it is not a practical algorithm. First, it is much more complex than the other algorithms, and second, its decisions depend on information farther in the future than the other algorithms. Nonetheless, its relative performance characteristics are of interest: we would like to understand whether or not the theoretical model we have defined actually gives insight into real system performance. If so, the theoretically near-optimal *reverse aggressive* can be used as a benchmark against which to compare other algorithms. The performance of *reverse aggressive* is our best *a priori* estimate of optimal performance.

The new algorithm *forestall* attempts to combine the best features of the other three algorithms: the good performance of *reverse aggressive* and the simplicity and implementability of *fixed horizon* and *aggressive*. *Forestall* avoids stalling while still making good (late) replacement decisions by estimating the point at which it needs to begin prefetching in order to prevent stalling.

1.4 Summary of results

In this paper we describe the results of a performance evaluation of the different policies for the d -disk integrated prefetching and caching problem. Our results from trace driven simulation demonstrate the practical performance characteristics of these algorithms. On our traces, we found that:

- All four algorithms significantly outperform demand fetching, even when advance knowledge of the access sequence is used to make optimal replacement decisions in conjunction with demand fetching.
- In compute-bound situations, *fixed horizon* and *forestall* have the best performance (which is usually matched by *reverse aggressive*'s).
- In I/O-bound situations, *aggressive* and *forestall* have the best performance (which is usually matched by *reverse aggressive*'s).

- In any given situation, one of *fixed horizon* or *aggressive* performs close to the theoretically near-optimal *reverse aggressive*.
- In all situations, *forestall* performs close to *reverse aggressive*.
- When data is well-laid out on the disks (e.g., striped), disk loads are balanced even without careful replacement choices. For this reason, *reverse aggressive* does not significantly outperform the other algorithms.
- *Fixed horizon* consistently places the least I/O load on the disks, due to its conservative fetching and near-optimal replacement choices. *Reverse aggressive* and *forestall* are intermediate between *aggressive* and *fixed horizon*.
- Batching of prefetch requests and disk head scheduling are crucial to the performance of prefetching and caching strategies.
- *Forestall* is a promising new approach that combines the best features of the other three algorithms: good performance regardless of I/O- or compute-boundedness, simplicity, and practicality.

We have focused on a rather narrow range of the input space: the single process, fully-hinted case. Clearly, prefetching and caching algorithms must deal effectively with missing or incorrect hints, as well as multiple simultaneously executing processes. *Fixed horizon*, *aggressive* and *forestall* can all be adapted to deal with these more general situations [5, 25].

1.5 Related work

Caching and prefetching have been known techniques to improve storage hierarchies for many years [2, 12]. In architectures, the work on caching and prefetching has focused on bridging the performance gap between CPU and main memory [28]. Research using caching and prefetching in database systems [9, 22, 10] showed that it is important to use applications' knowledge to perform caching and prefetching.

File caching and prefetching have become standard techniques for sequential file systems [12, 19, 14, 21, 29, 4, 13, 7, 25]. The most common prefetching approach is to perform sequential readahead [12, 19, 20]. The limitation of this approach is that it only benefits applications that make sequential references to large files. Another large body of work has been on predicting future access patterns [11, 29, 22, 10, 13]. Recently, caching and prefetching have also been studied for parallel file systems [11, 17, 24].

Although much work has been done in file caching and prefetching, most of it has considered one or the other in isolation. Recent studies for the single disk case showed [6, 5] that it is important to integrate prefetching, caching and disk scheduling together and that a properly integrated strategy can perform much better than a naive strategy. For the multi-disk case, a theoretical study [16] presented and analyzed *aggressive* and *reverse aggressive*. Other parallel prefetching strategies include one stripe lookahead prefetching on RAID

arrays, and Patterson *et al.* [25]’s TIP2 system. The one stripe lookahead benefits only applications that use large files, and would perform little prefetching for other applications. TIP2 uses the *fixed horizon* algorithm we have studied here. Patterson *et al.* [25] also present a cost-benefit technique for controlling buffer allocations for both hinting and non-hinting applications in a multi-process environment.

Previous studies of the algorithms considered here have been incomparable. Differences in hardware, both in the processor and the I/O system, as well as in the benchmarks used to evaluate the algorithms, have made it difficult to understand the differences between them. This paper represents the first direct comparison of these approaches. Using the knowledge learned from this comparison, we have designed a new algorithm that attempts to combine the best features of the previous efforts.

1.6 Organization of the paper

In the next section we describe the first three algorithms and their theoretical basis. In section 3, we describe our simulation framework. In section 4, we present the results of our simulations using the first three algorithms. In section 5 we describe the new algorithm *forestall* and present simulation results on its performance. We present our conclusions in section 6.

2 The algorithms

We begin by introducing the framework used to study this problem and the terminology used in the rest of the paper.

2.1 Theoretical model

Our theoretical model consists of two levels of memory hierarchy: a cache of K data blocks, and d (disk) storage devices. The execution of a program makes a known *request sequence* of references r_1, r_2, \dots, r_n to a set of m data blocks.

If a reference hits in the cache, it takes one time unit. Otherwise, the missed block must be fetched from a storage device. The system can either fetch a block on a miss (demand driven) or fetch the block before it is referenced in anticipation of a miss (prefetch). Either case takes F time units. If the cache is full, a cache block must be *evicted* before the fetch is issued to make room for the requested data block. While the fetch is in progress, neither the incoming block nor the discarded block is available for access.

We assume that each block resides on a single disk. Fetches to a single disk are serialized, but fetches on different disks can be executed concurrently.

When the program tries to access a block that is not in the cache, it stalls until the block arrives in the cache. The stall time is either F if the block is fetched on demand or $F - i$ if the block is prefetched i time units before the reference. The measure of performance is the elapsed time required to serve the entire request sequence; this is equal to the number of references plus the total stall time.

The goal is to minimize application elapsed time, by deciding when to fetch a block from a disk, which disk to fetch from, which block to fetch, and which cache block to evict (when the cache is full).

The time unit models the CPU time spent between two consecutive file references — the CPU time includes the time to copy the accessed file data from kernel address space to a user address space buffer, and the time for the application to consume the file data. The model simplifies the real situation by assuming that the CPU time between every two file references is the same, that all disk accesses take the same amount of time, and that there is no CPU overhead incurred by issuing an I/O request. These simplifications were made in order to make the problem theoretically tractable. Our simulations use actual CPU times collected in our traces and an accurate simulation model of modern disk drives, and charge a driver overhead for each request made to a disk.

2.2 Optimal prefetching rules

The following simple rules can be assumed of any optimal strategy in the single-disk case [6].²

- *Optimal fetching*: when fetching, always fetch the missing block that will be referenced soonest;
- *Optimal replacement*: when fetching, always evict the block in the cache whose next reference is furthest in the future;
- *Do no harm*: never evict block A to fetch block B when A ’s next reference is before B ’s next reference;
- *First opportunity*: never evict A to fetch B when the same thing could have been done one time unit earlier.

Unfortunately, as exhibited in the example in section 1, some of these rules no longer hold in the multiple-disk case. It may be necessary to violate all of the rules except *first opportunity* to produce an optimal schedule.

2.3 The *fixed horizon* algorithm

As described earlier, the *fixed horizon* algorithm is based on the TIP2 system running a single hinting process [25].

Fixed horizon: Whenever there is a missing block at most H references in the future, issue a fetch for that block, replacing the cached block whose next reference is furthest in the future, provided that reference is further than H accesses in the future (which will certainly hold if $H < K$).

Fixed horizon is consistent with the first three rules of optimal prefetching for a single disk. An advantage of not following the fourth rule is that *fixed horizon* needs less information about references beyond the prefetch horizon than the other algorithms. A disadvantage is that when additional

²These rules are optimal in the sense that any schedule that does not follow them can be transformed into one that does, with performance at least as good.

information is available, *fixed horizon* can have elapsed time nearly twice optimal.

The prefetch horizon H is computed as the ratio of the average time it takes to read a block from disk and the minimum time it takes to consume a single block of data. In the theoretical model, $H = F$.

Fixed horizon tries to fetch as late as possible without stalling in order to make the best possible replacement decision. Each fetch is issued so that it will complete just in time for the reference. If parallelism increases to the point that each request is made to an idle disk, this algorithm is optimal. However, in practice, a sufficient number of disks may not be available. In this case, *fixed horizon* may initiate fetches too late to avoid stalling. In fact, because it never initiates a fetch more than H references ahead of the missing block, *fixed horizon* may allow a disk to become idle even though the future requests beyond the prefetch horizon contain many missing blocks. On the other hand, if the missing blocks in the sequence tend to be separated by many intervening references to blocks that are present in the cache, we'd expect *fixed horizon* to have performance much closer to optimal than its worst case.

2.4 The aggressive algorithm

The (multi-disk) aggressive algorithm is based on the Cao *et al.* (single-disk) aggressive algorithm [6], which is provably near-optimal in the single-disk case.

(Multi-disk) aggressive: Whenever a disk is free, prefetch the first missing block on that disk, replacing the block whose next reference is furthest in the future, under the condition that the next access to the evicted block is after the next access to the block being fetched.

Aggressive is the most aggressive prefetching strategy that is consistent with the four optimal prefetching rules described in section 2.2. As mentioned, some of these rules are no longer valid in the multiple disk case. This provides some of the intuition for the following theorem.

Theorem 1 [16] *For any access pattern, and any layout of data on disks, the elapsed time of aggressive is at most $d(1+\epsilon_1)$ times that of the optimal elapsed time (the minimum possible), where d is the number of disks, and ϵ_1 is a small constant that depends on system parameters.*³

There are worst case access patterns/data layouts for which the elapsed time of aggressive is at least d times the minimum possible.

It is important to note that this worst case result depends on access patterns and data layouts in which the load is heavily unbalanced between the disks. If the request sequence is balanced, *aggressive* has near-optimal performance.

³ ϵ_1 here is F/K where F is the fetch time/compute time ratio and K is the cache size measured in blocks. For typical system parameters ϵ_1 is less than 0.02.

2.5 The reverse aggressive algorithm

The *reverse aggressive* algorithm exploits global knowledge in order to produce a prefetching schedule that achieves near-optimal elapsed time. It does this by balancing disk workload through carefully selected replacement decisions.

Reverse aggressive: Construct a prefetching schedule for the *reversed* sequence that *replaces* at most one block on each disk in parallel as follows: Whenever a disk is free, determine the block B not needed for the longest time on that disk. If the next request to B is after the first missing block, issue a fetch for the missing block, *replacing* B . Transform this prefetching schedule back to a schedule for the original sequence by treating each fetch on the reverse sequence as an eviction on the forward sequence and vice versa.

For a proof of correctness, more details on how and why this algorithm works well, and a proof of the following theorem, see [16].

Theorem 2 [16] *For any request sequence, and for any layout of the data on the disks, the elapsed time of reverse aggressive is at most $1 + \epsilon_2$ times the optimal elapsed time.*⁴

There are two key properties of *reverse aggressive* that result in this theorem. First, whereas *aggressive* chooses evictions without considering the relative loads on the disks, *reverse aggressive* greedily evicts to as many disks as possible on the reverse sequence. In the forward direction, this translates to performing a maximal set of fetches in parallel. The fact that these are fetches in the forward direction means that at some point earlier in the sequence, corresponding blocks were evicted. Thus the eviction decisions of *reverse aggressive* on the forward sequence are based on the ability to prefetch the evicted blocks later on in parallel. Second, whereas *aggressive* can wastefully prefetch ahead on some of its disks, *reverse aggressive* is greedy in the reverse direction. Consequently, it is fetching blocks in the forward direction just in time (to the extent possible) for them to be used. This results in performing close to the best evictions possible for those fetches.

2.6 Practical considerations

Several important features of real systems are not captured by our theoretical model.

1. Disk response times and CPU times between I/O requests are not constant.

We use average values for each and expect that variation in event times does not substantially invalidate the algorithm's decisions. In our experimentation, this does not appear to be a major effect, with one exception (see section 4.3). (The systematic effects of disk scheduling on disk response time are considered separately).

⁴ ϵ_2 here is less than Fd/K , where F is the fetch time/compute time ratio, d is the number of disks, and K cache size in blocks. For typical system parameters, ϵ_2 is less than 0.1, and sometimes significantly less.

2. Access patterns exhibit locality of reference, and loads are balanced across the multiple disks when data is striped.

In practice, this allows both *fixed horizon* and *aggressive* to effectively utilize multiple disks, and achieve elapsed times comparable to the theoretically superior *reverse aggressive*.

3. Disk accesses require significant CPU overhead to form the request, communicate with the disk, and service the resulting interrupt(s). Thus, avoidable data fetches may add elapsed time even if they do not cause stalls.

Because the theory assumes that fetches entail no CPU overhead, this penalty punishes overly aggressive fetching. In practice, this effect favors the *fixed horizon* algorithm since its late replacement decisions tend to lead to the fewest fetches.

4. Disk response time is sensitive to the order in which requests are serviced.

In particular, disk scheduling reduces average disk response time as more accesses are presented and allowed to be reordered by the disk (driver). Although *fixed horizon* implicitly allows multiple outstanding requests at each disk, *aggressive* and *reverse aggressive* were defined to submit only one request at a time, since in the theoretical model there is no advantage to batching. Because of the significance of the disk scheduling effect, we modify the definitions of *aggressive* and *reverse aggressive* to submit disk requests in batches. We have found that the performance of all three algorithms benefits from the CSCAN disk scheduling algorithm.

Reverse aggressive also benefits from batching of requests during its construction of its prefetching schedule (the reverse pass over the request sequence). This is because typical request sequences exhibit spatial locality; by batching requests on the reverse pass, *reverse aggressive* generates missing blocks to be fetched on the forward sequence in groups that exhibit locality of reference.

The inter-request CPU time is actually composed of two components, a fixed amount of time to read a block out of the cache, and a variable amount of time to process the data. Our implementation of *fixed horizon* assumes the data processing time to be zero, and uses the ratio of the average disk response time to the time to read a block from the cache as the prefetch horizon H . This ensures that any prefetch to an idle disk will complete in time for the reference. Assuming an average disk response time of $15ms$ (which is usually an overestimate in our simulations) and $243\mu s$ to read a block from the cache (which was measured on the implemented TIP2 system) yields a value of $H = 62$; we used this value in all our simulations, except where noted otherwise.

2.7 Implementations of the algorithms

In the context of the considerations of the previous section, we summarize the implementations we compared.

Fixed horizon: Whenever there is a missing block at most H references away, issue a fetch for that block, replacing the block whose next reference is furthest in the future. Note that this algorithm may at any time have up to H outstanding references to a disk yielding opportunities for disk scheduling.

Aggressive: Whenever a disk D is free, construct a batch of at most `batch-size`⁵ fetches to initiate on D as follows: As long as the first missing block B on disk D precedes the block B' whose next request is furthest in the future, add the fetch/eviction pair B/B' to the batch. Issue the batch.

If two or more disks are free at the same time, we consider all their missing blocks together, in order of increasing request index. Each next missing block is issued to the appropriate disk (and the best possible choice of evictions is made), if the disk's batch is not full and the *do no harm* rule allows it. At some point, either the last free disk's batch becomes full or the *do no harm* rule disallows issuing further requests.

Reverse aggressive: Assuming a fixed ratio F between the time for a disk access and the inter-reference CPU time, consider the reversed sequence, and use it to derive a prefetching schedule as described in section 2.5, but construct the schedule in batches as done by *aggressive*.

This prefetching schedule is then transformed into a schedule of fetch/eviction pairs for the forward sequence. Associated with each eviction is a *release time*, the earliest index in the request sequence at which the block can be evicted (i.e. one greater than the index of the last request to the block until it is possibly fetched back into the cache at some later time.) The eviction choices are naturally ordered by increasing release point due to the method used by *reverse aggressive* to construct its schedule. Fetches may need to be re-ordered according to increasing request index; they are then matched to eviction choices according to these orderings.

This schedule is used to drive the disk model as follows. Whenever a disk D is free, add the first up to `batch-size` fetch/eviction pairs B_i/B'_i that have been released, and for which B_i resides on disk D , to the batch. Issue the batch.⁶

Notice that *aggressive* and *fixed horizon* use less lookahead information than *reverse aggressive*, in that for both of them, the “only” future information needed are the identities of the next missing blocks (up to H missing blocks for *fixed horizon*, and up to d times `batch-size` for *aggressive*), and their positions in the sequence relative to the next references to blocks currently in the cache.

3 Simulation framework

We used trace-driven simulation to evaluate the performance of the algorithms. We believe our simulation model to be an accurate reflection of the practical performance characteristics of the algorithms. The reference streams are taken from traces of real applications' behavior; the trace information

⁵The batch sizes used are listed in table 6.

⁶The batch sizes and estimate F used by *reverse aggressive* are discussed in section 4.4.

Sector size	sectors per track	tracks per cylinder
512 bytes	72	19
cylinders	rotational speed	disk cache size
1962	4002 rpm	128 Kbytes
ave. access time (8Kbyte)	controller interface	transfer rate
22.8ms	SCSI-II	10 MB/sec

Table 1: HP 97560 characteristics.

xds elapsed times (secs)				
	CMU simulator		UW simulator	
disks	F.H.	Agg.	F.H.	Agg.
1	63.3	61.6	65.6	63.7
2	36.9	34.1	38.0	34.3
3	33.6	33.9	36.2	33.7
4	33.8	35.1	34.2	35.1
5	33.0	34.2	33.5	34.4
synth elapsed times (secs)				
	CMU simulator		UW simulator	
disks	F.H.	Agg.	F.H.	Agg.
1	213.0	168.5	201.4	155.8
2	136.3	126.9	130.9	121.7
3	118.9	149.5	118.9	150.4
4	118.9	150.4	118.9	150.1

Table 2: Comparison of the simulators on the xds and synth traces.

we use is unaffected by prefetching and caching activity. The accurate modelling of disk fetch times, I/O driver overhead costs, and application process compute times in the simulations is a key difference relative to the theoretical framework. However, our simulators do not model serialization of DMA transactions.⁷

Two separate simulators were developed, one at Washington (UW) and one at Carnegie Mellon (CMU). The UW simulator uses the disk drive simulation of Kotz *et al.* [18] (which is based on that of Ruemmler and Wilkes [26]) to accurately model I/O costs. This simulation models fine architectural details to provide a very accurate simulation of the HP 97560 disk drive. Table 1 lists several characteristics of the HP 97560 (taken from [26]). The CMU simulator uses the Berkeley RaidSim [8] simulator, as modified at CMU, to simulate 0661 IBM Lightning disk drives.

The simulators were cross-validated on a common set of traces. The CMU simulator does not implement *reverse aggressive*. We obtained good agreement between the simulators on the results for *aggressive* and *fixed horizon* for several traces. Table 2 shows the elapsed times measured by the simulators for the xds and synth traces described below. Remaining differences between the simulators are consistent with the differences in the disk models. We report here results for all algorithms obtained using the UW simulator.

In our simulations, we ignore write operations. Write performance is less critical to I/O performance since the application generally does not have to wait for the disk to be

⁷We do not expect this to have a significant effect on the results since the DMA time is much less than the disk access time.

trace	reads	distinct blocks	compute time (sec)
dinero	8867	986	103.5
cscope1	8673	1073	24.9
cscope2	20206	2462	37.1
cscope3	30200	3910	74.1
glimpse	27981	5247	38.7
ld	5881	2882	8.2
postgres-join	8896	3793	11.5
postgres-select	5044	3085	79.2
xds	10435	5392	30.8
synth	100000	2000	99.9

Table 3: Trace summary data.

written. Moreover, the impact this has on the results is small since most of the references in our traces are reads.

We simulated disk arrays of sizes 1-8, 10, 12, and 16. Most of our figures show a smaller range of sizes, however. In each case, the performance with a larger number of disks is the same as that with the largest number of disks shown.

3.1 File access traces

We used a set of traces collected on a DECstation 5000/200. The running time of all the applications is dominated by disk read accesses. Each trace consists of a sequence of file block read requests in the order they were issued, and the sequence of measured process compute times between read requests, of a single execution thread. We used an I/O driver overhead of .5ms per I/O operation, which is typical of the 5000/200.

The applications are:

cscope[1-3]: an interactive C-source examination tool written by Joe Steffen, searching for eight symbols (cscope1) in a 18MB software package, searching for four text strings (cscope2) in the same 18MB software package, and searching for four text strings (cscope3) on a 10MB software package. With multiple queries, cscope will read multiple files sequentially multiple times.

dinero: a cache simulator written by Mark Hill. This application reads one file sequentially multiple times.

glimpse: a text information retrieval system from the University of Arizona, searching for four keywords in a 40MB snapshot of news articles. It builds approximate indexes for words to allow both relatively fast search and small index files. The result is that the index files are accessed repeatedly, whereas the data files are accessed infrequently.

postgres-join: the Postgres relational database system developed at the University of California at Berkeley, performing a join between an indexed 32MB relation and a non-indexed 3.2MB relation. The relations are those used in the Wisconsin Benchmark [3]. Since the result relation is small, most of the file accesses are reads. Here, the index blocks are accessed much more frequently than the data blocks.




Figure 2: Performance on the postgres-select trace. Each group of bars represents the performance of the four algorithms *optimal demand fetching*, *fixed horizon*, *aggressive*, and *reverse aggressive*, in left-to-right order.

postgres-select: the Postgres relational database system executing a selection query of choosing 2% of the tuples from an indexed 32MB relation. The selection query is part of the Wisconsin Benchmark suite [3] and uses indexed search.

ld: the Ultrix link-editor, building the Ultrix 4.3 kernel from about 25MB of object files.

xds: a 3-D data visualization program, XDataSlice, generating 25 planar slice images at random orientations from a 64MB data file.

Finally, we used a synthetic trace **synth** containing 50 passes through a loop of 2000 sequential blocks. Compute times between read requests were generated according to a Poisson distribution with a 1 ms mean.

Table 3 shows the length (number of read requests), number of distinct blocks requested, and total application compute times for each of the traces.

The cache size was set to be 10MB (or $K = 1280$ blocks of 8 kbytes each) for all traces except *dinero* and *cscope1*. These traces contain references to fewer than 1280 distinct blocks. For these traces, the cache size was reduced to 4MB (512 blocks). We assume the cache to be empty (or to contain some other application’s data) when the traced application starts. The entire cache is available to the traced application.

3.2 Data placement and disk head scheduling

The data was striped across the array using a one-block stripe unit. Some of our traces represented block numbers by (file,offset) pairs; for these we chose a random starting point within a group of 8550 8kbyte blocks (which occupy 100 cylinders on the HP 97560) for each file, corresponding to typical file system clustering mechanisms. The maximum seek time within a group of 100 cylinders is 7.24ms. Thus, in our simulations the average response time is typically lower than the 22.8ms listed in table 1. Other traces referred to logical filesystem block numbers; for these traces we used the actual block number for each access. Except where noted, we use CSCAN disk head scheduling.

4 Results

In the following sections, we examine the behaviors of the algorithms in detail. We begin by comparing the performance of the algorithms with that of demand fetching. We then examine the algorithms’ performance on the synthetic trace, an easily understood access pattern that illustrates the key differences in behavior between the algorithms. Next we examine performance on the application traces, and explore the effects on the results of changes in various simulation parameters.

4.1 Comparison with demand fetching

In order to make this comparison as favorable as possible to demand fetching, we use the optimal offline replacement policy: whenever a block is fetched, the block in the cache whose next reference is furthest in the future is replaced. Figure 2 shows the elapsed times of the three algorithms and of optimal demand fetching on the postgres-select trace for varying numbers of disks between one and sixteen. The elapsed times are divided into three components: process compute time, I/O driver overhead (processor) time, and the time the processor spends idle, stalling on I/O. From this figure we see that (1) all three prefetching algorithms significantly outperform optimal demand fetching, and (2) the three prefetching algorithms achieve near linear reduction in I/O overhead until the applications become compute-bound. These two behaviors are consistent across all the applications we have studied.

4.2 Fundamental differences

The synthetic trace is used to examine the algorithms’ behavior on a simple, known sequence in order to gain insight into the algorithms’ performance. This trace shows the relative behaviors typical of the three algorithms in exaggerated form. Figure 3 summarizes the results for one to four disks.

The sequential accesses allow excellent performance from the disks; average response times are between 3 and 4 ms. In each case, *fixed horizon* performs 38000 fetches, 720 more




Figure 3: Performance on the synth (left) and cscope1 (right) traces. Each group of bars represents the performance of the three algorithms *fixed horizon*, *aggressive*, and *reverse aggressive*, in left-to-right order.

than the minimum possible 37280 performed by optimal demand fetching. (The total sequence length is 100,000).

With a single disk, the synthetic application is I/O bound. *Fixed horizon*'s conservative prefetching strategy reduces I/O stalling relative to demand fetching, but not as much as *aggressive*'s and *reverse aggressive*'s more aggressive strategies. After each pass through the loop under *fixed horizon*, the cache contains 1280 sequential blocks and the other 720 blocks in the sequence are not cached. The clustering of the 720 missing blocks allows good disk performance; however, the clustering of the 1280 cached blocks causes *fixed horizon* to leave the disk idle until the last H cached blocks are being read. *Aggressive* and *reverse aggressive* perform 39240 and 39265 fetches, respectively, slightly more fetches than *fixed horizon*'s 38000, resulting in a small difference in driver overhead. However, they are able to eliminate much of the I/O stall time by prefetching distant blocks and thus not idling the disk appreciably.

With two disks, *fixed horizon* is able to eliminate most of the stall time, without increasing the total number of fetches. *Aggressive* has nearly eliminated stall time completely, but at a higher driver cost due to its increased number (41902) of fetches. *Reverse aggressive* is between *fixed horizon* and *aggressive* in stall time; it performs 42000 fetches. Elapsed times are similar under all three algorithms. This case marks the transition from I/O-boundedness to compute-boundedness.

With three disks, stall time has been eliminated completely by all three algorithms. *Aggressive* uses the excess I/O bandwidth to prefetch and subsequently evict every block for every reference. In fact, because *aggressive* is willing to prefetch significantly ahead on one disk relative to others, it wastes 994 fetches, replacing a prefetched block from the cache before it is used in order to fetch a block on a different disk that will be needed sooner. Fortunately, this effect does not increase as the number of disks increases since with increasing I/O bandwidth, *aggressive*'s prefetching becomes so successful that every fetch is to the first missing block in the future. Such a block can never be replaced before it is used, since that would violate the *do-no-harm* rule.

Nonetheless, the elimination of stall time by *aggressive* comes at a high cost: the driver overhead for the extra fetches

pushes *aggressive*'s elapsed time higher than the two-disk case. In contrast, *fixed horizon* prefetches far enough ahead to serve all requests without stall, but no farther. Dedicating at most H buffers to prefetching, *fixed horizon* is able to eliminate stalling altogether without any additional fetches. *Reverse aggressive* performs 37907 fetches, fewer than *fixed horizon*, also eliminating stall time.

4.3 Application traces

The application traces show differences among the three algorithms similar to those shown by the synthetic trace, but less pronounced.

The right portion of figure 3 shows the performance of the three algorithms on the CPU-bound cscope1 trace. The behavior here is similar to that for the synthetic trace: *aggressive* eliminates stalling but issues too many fetches resulting in a greater driver overhead.

At the I/O-bound end of the spectrum, figure 4 shows a detailed breakdown of the performance of the three algorithms on the ld trace, from one to sixteen disks. With one disk, all three algorithms are I/O bound and have comparable performance. From two to eight disks, the more aggressive prefetching of *aggressive* and *reverse aggressive* results in somewhat less stalling than *fixed horizon*. At ten disks, *fixed horizon*'s performance matches *aggressive*'s. Beyond this point, the tradeoff between excessive stalling caused by leaving disks idle, and excessive driver overhead caused by prefetching aggressively, favors *fixed horizon* over *aggressive*. The other traces reflect similar trends, with different points of crossover: above five disks for postgres-select, glimpse, and cscope2, and below five disks for postgres-join, dinero, cscope1, and xds.

An exception to the generally best performance of *reverse aggressive* is the cscope3 trace, shown in figure 5. Note that *reverse aggressive*'s performance is much worse than *aggressive*'s with one disk. This is a case in which the differences between the theoretical model and the simulation model affect the performance of *reverse aggressive*. Recall that since *reverse aggressive* is offline, it generates a complete schedule based on its *estimate* of F . When it uses a smaller estimate of F , each fetch is assumed to complete earlier (relative to




Figure 4: Performance on the ld trace.

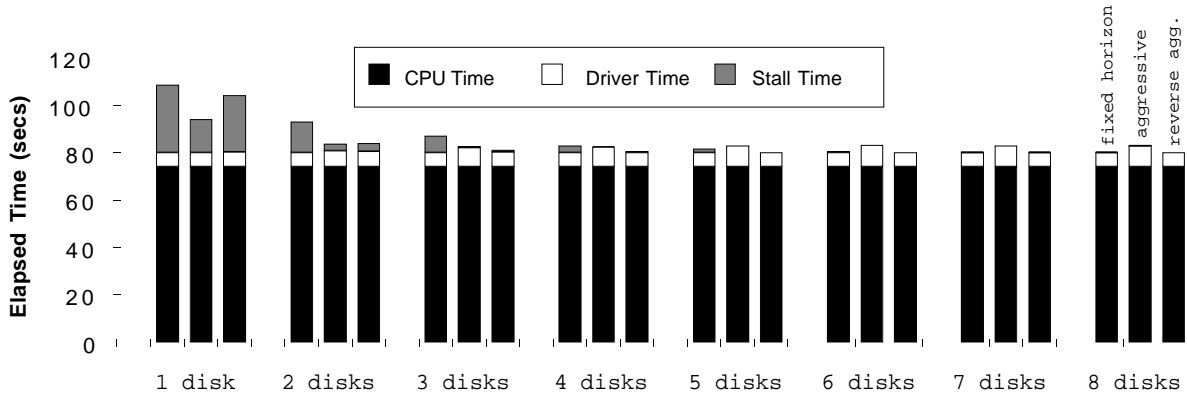


Figure 5: Performance on the cscope3 trace.

the inter-reference compute time) and therefore *reverse aggressive* generates a more aggressive prefetching schedule that keeps the disk(s) busier. When it uses a larger estimate of F , each fetch is assumed to take longer, and therefore *reverse aggressive* must delay the scheduling of subsequent fetches in the sequence, thus generating a more conservative prefetching schedule. In our implementation of *reverse aggressive*, the single best estimate of F is used for each trace. On traces with large variation in inter-reference compute times, any single estimate of F will be either too small or too large for some parts of the trace. This is the case for *cscope3* – examination of the trace reveals that the inter-reference compute times are bursty. Runs of compute times near 1ms are interspersed with runs of times around 7ms. Since the average fetch time on this trace with one disk is about 8ms, the ratio of fetch time to compute time (the “true” value of F) varies from about 1 to about 8.

In fact, with a single disk, *aggressive* has the same theoretical performance bounds as *reverse aggressive*. It is not surprising that *aggressive*’s inherent adaptivity to varying fetch times and compute times should give it an advantage over *reverse aggressive* in this case. This effect is noticeable, but less pronounced, on the *synth* trace as well.

On the remaining traces, *reverse aggressive*’s elapsed time varies from 3.6% worse to 10.7% better than the superior of *fixed horizon* and *aggressive* in any given configuration. For the full data, see the appendix.

disks	demand fetching	<i>fixed horizon</i>	<i>aggressive</i>	<i>reverse aggressive</i>
1	.82	.98	.99	.98
2	.41	.90	.92	.92
3	.27	.82	.87	.85
4	.20	.72	.81	.80
5	.16	.66	.70	.69
6	.13	.58	.63	.60
7	.12	.50	.62	.50
8	.10	.45	.56	.42
10	.08	.36	.43	.35
12	.07	.30	.35	.28
16	.05	.22	.26	.21

Table 4: Disk utilization on the postgres-select trace.

Table 4 shows the utilization of the disks (averaged over the disks when there are more than one) for demand fetching and the three prefetching algorithms on the postgres-select trace. For moderate numbers of disks, *aggressive* places the greatest load on the disks, followed by *reverse aggressive* and then *fixed horizon*; demand fetching places the least load on the disks. With a very high degree of disk parallelism, *reverse aggressive*’s offline schedule places even less load on the disks than *fixed horizon*’s conservative strategy.

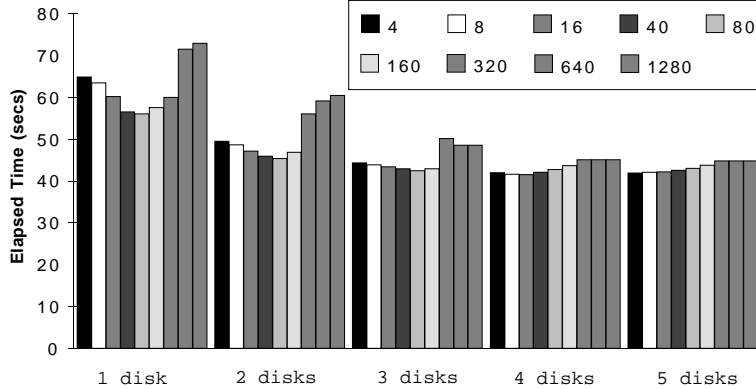


Figure 6: Performance of *aggressive* on the *cscope2* trace, as a function of the batch size.

4.4 Varying parameters

The performance of the algorithms depends on a set of parameters which interact in complicated ways with the applications' access patterns and inter-reference compute times, the layout of data on disks, the disk-scheduling discipline, and the characteristics of the disks. In this section, we explore the behavior of the algorithms when some of these parameters are varied. In this section, we present general observations and only a small portion of the data. For the full data, see the appendix.

We have already described most of the primary effects that explain what we see. These are:

- *scheduling*: an increase in the number of outstanding fetches issued by a prefetching algorithm results in increased latitude to reorder fetches and thus reduced disk response times. This effect is strongest in I/O-bound situations.
- *out-of-order fetching*: reordering of fetches can increase stall penalties when early missing blocks are fetched after later missing blocks. This effect is strongest in CPU-bound situations where any stall penalty is costly. When there is significant stalling, this effect is masked by other stalls and compensated for by the reduced average response time.
- *early replacement*: as prefetching becomes more aggressive, inferior replacement choices are made, leading to more fetches and in many cases, an increase in elapsed time.
- *limited aggressiveness*: the extent to which an algorithm can prefetch is limited by the *do no harm* rule.

Disk-head scheduling

The results shown in the previous section were obtained using CSCAN disk-head scheduling. CSCAN was used rather than SCAN since the HP 97560 contains a read-ahead buffer; CSCAN always scans in the same direction that the disk reads, improving the hit rate in the read-ahead buffer. We compared the performance impact of CSCAN

disks	<i>fixed horizon</i>	<i>aggressive</i>	<i>reverse aggressive</i>
1	14.9	19.2	24.0
2	4.85	11.3	22.1
3	2.59	8.36	19.9
4	0.53	3.59	6.71
5	-0.62	-0.77	0.0
6	-0.68	-0.31	0.0
7	-2.15	-0.45	0.0
8	-0.42	-0.17	0.0
10	-0.05	0.09	0.0
12	0.0	0.11	0.0
16	0.0	0.0	0.0

Table 5: Percentage improvement of CSCAN over FCFS on the *postgres-select* trace.

disk-head scheduling versus FCFS scheduling. Relative to FCFS, CSCAN improves the performance of *reverse aggressive* the most, up to 24%, and that of *fixed horizon* the least, up to 15%. For *aggressive*, the greatest benefit was 19%. Because of out-of-order fetching, CSCAN sometimes degrades performance slightly relative to FCFS in compute-bound situations. This effect is strongest for *fixed horizon* since it issues fetches later than they are issued by the other algorithms. The maximum degradation we observed is 3.6% (for *fixed horizon* with six disks on the *glimpse* trace).

Table 5 shows the performance benefit of CSCAN scheduling relative to FCFS on the *postgres-select* trace for all three algorithms with 1-16 disks.

The batch size used by *aggressive*

Figure 6 shows the effect of varying *aggressive*'s batch size on the *cscope2* trace. For each number of disks, performance initially improves with increasing batch size due to improved scheduling. For example, for one disk, the average fetch time drops from 10.4ms to 8.4ms as the batch size increases from 4 to 160. Eventually, out-of-order fetching and early replacement become more important and performance drops off again. For example, for one disk the number of fetches increases from 6771 to 9806 as the batch size increases from 160 to 1280.




Figure 7: Performance of *fixed horizon* as a function of the prefetch horizon H on the *cscope1* (left) and *cscope2* (right) traces.

1 disk	2 disks	3 disks	4 disks
80	40	40	16
5 disks	6 disks	7 disks	> 7 disks
16	8	8	4

Table 6: Batch sizes used for *aggressive*.

As the number of disks increases, the variation in performance with batch size diminishes, and the best batch size shifts to a smaller value. This is because in more compute-bound situations, out-of-order fetching and limited aggressiveness are the dominant factors. Because of limited aggressiveness, the number of fetches increases only from 11325 to 11399 as batch size increases from 160 to 1280 with 5 disks.

Although the optimal batch size decreases with the number of disks for all the traces, it varies significantly from trace to trace. For example, for the *xds* trace, the optimal batch size for one to three disks was 16, and for four or more was 4. All the results for *aggressive* presented in section 4.3 were obtained using the batch sizes given in table 6. The performance of *aggressive* with these fixed batch sizes is on average 0.7 % worse (and at most 11% worse) than its performance with the best batch size for the configuration.

Prefetch horizon

The left side of figure 7 shows the effect of varying *fixed horizon*'s prefetch horizon H on the *cscope1* trace. We see that for each number of disks, performance deteriorates with increasing H (except on one disk, where it improves slightly until $H = 64$ is reached). This is due to out-of-order fetching and early replacement. For example, with 1 disk, earlier replacements cause the number of fetches to increase from 4959 with $H = 64$ to 8535 with $H = 2048$. Out-of-order fetching accounts for all the stall time with 2 and 3 disks when $H \geq 512$; using FCFS scheduling this stall time is eliminated.

On the more I/O bound traces such as *cscope2*, also shown in figure 7, we find a significant initial performance improvement with increasing H because the more aggressive prefetching eliminates stalling. Only at very large values of H does performance decline again.

The parameters used by *reverse aggressive*

We experimented with the batch size and fixed value of F used by *reverse aggressive* to construct its schedule on its reverse pass over the request sequence, as well as the batch size used on the forward pass. Since we use *reverse aggressive* only as a benchmark against which to compare the other algorithms, the main purpose of these experiments was to determine the optimal configuration (choice of F and batch sizes) for each trace, for each number of disks.

These experiments show that, as with *aggressive*, a smaller (resp. larger) batch size benefits a more compute-bound (resp. I/O-bound) application. Recalling that as *reverse aggressive*'s estimate of F decreases, it becomes increasingly aggressive, we similarly find that a smaller (resp. larger) value of F benefits a more I/O-bound (resp. compute-bound) application.

Processor speed and cache size

In order to assess the impact of improved CPU performance relative to disk performance, we ran our trace-driven simulations assuming a processor twice as fast. For these tests, *fixed horizon*'s prefetch horizon H was doubled to 124. The results are entirely unsurprising: faster processors are more dependent on I/O performance so that the payoff of using multiple disks and prefetching is increased. In addition, since a larger number of disks is needed to eliminate I/O overhead, the point at which the tradeoffs begin to favor *fixed horizon* over *aggressive* is shifted to a larger number of disks. This behavior was consistent across the applications.

In order to assess the impact of cache size on performance, we ran our trace-driven simulations with varying cache sizes: 640, 1280, and 1920 blocks. As cache size increases, the performance of all the algorithms improves. In I/O-bound cases, a larger cache improves *aggressive*'s and *reverse aggressive*'s performance more than *fixed horizon*'s since they prefetch more aggressively. In more compute-bound cases, *aggressive*'s excessive driver overhead penalizes it even more with a larger cache, so that *fixed horizon*'s performance relative to *aggressive* improves slightly as cache size increases. This is illustrated in table 7, which shows the performance of *fixed horizon* relative to *aggressive* as percentage difference, as a

cache size	1 disk	2 disks	4 disks	8 disks	16 disks
640	6.0	14.7	24.8	7.3	-2.6
1280	11.3	20.2	24.5	5.7	-3.8
1920	13.8	25.0	21.7	5.7	-3.8

Table 7: Elapsed time as a function of the cache size and number of disks of *fixed horizon* relative to *aggressive* (percentage difference) on the glimpse trace.

function of the cache size and the number of disks on the glimpse trace.

5 A new approach

We have designed a new algorithm, *forestall*, attempting to combine the best features of all three previously described algorithms: the good performance of *reverse aggressive* regardless of I/O-boundedness or compute-boundedness, and the simplicity and implementability of *fixed horizon* and *aggressive*. *Forestall* tries to avoid stalling while still making good (late) replacement decisions by estimating the point at which it needs to begin prefetching in order to prevent stalling. It makes this estimate based on its current cache state.

Returning to the theoretical model, suppose that there is a single disk, and that at some point during the servicing of the request sequence, the cache contains the next $2F - 1$ blocks requested. (Recall that in the theoretical model, the interreference CPU time is taken to be 1 time unit, and the time to fetch a block from disk is F time units.) Further suppose that the subsequent two requests are missing from the cache. *Aggressive* immediately starts fetching and avoids stalling on the missing blocks, bringing the second missing block into the cache at time $2F - 1$ just in time to serve the request without stalling. *Fixed horizon* leaves its disk idle until the cursor is within F requests of the first missing block; it stalls $F - 1$ steps on the second missing block. In contrast, suppose there is only one missing block at a distance of $2F - 1$ from the cursor. In this case, *aggressive* will fetch immediately and make a possibly inferior replacement choice. *Fixed horizon* waits until its cursor is within F steps of the missing block, and prefetches just early enough to avoid stalling; in the intervening time, it may have finished using a block that is not needed until later in the sequence (if at all) than the one evicted from the cache by *aggressive*.

Forestall behaves as does *aggressive* in the first case, and as does *fixed horizon* in the second. For each i , $i \geq 1$, let d_i denote the distance from the cursor to the i^{th} missing block in the request sequence. For any $i \geq 1$, if $iF > d_i$, processing will surely stall on the i^{th} missing block or some earlier missing block. It will take iF time units to fetch the first i missing blocks, and at most the next d_i requests can be served concurrently. *Forestall* initiates a prefetch according to the *optimal fetching* and *optimal replacement* rules whenever $iF \geq d_i$ is true for some i and the *do no harm* rule allows it.

Practical considerations

As do the other algorithms, *forestall* requires modifications in order to account for differences between the theoretical model and real systems. Requests need to be issued in batches in order to reduce average disk access times. The ratio F of disk response time to interaccess time is not constant and must be estimated. In our implementation, we estimate F by tracking recent disk response times and compute times: F is dynamically computed on a per-disk basis as the ratio between the sum of the most recent 100 disk access times and the most recent 100 interreference CPU times.

Just as we needed the prefetch horizon H to be an overestimate of F for *fixed horizon* to have adequate performance, *forestall's* performance depends on overestimating F in certain situations as well. We denote by F' the overestimate of F used by *forestall*. We evaluated *forestall's* performance with different values of the parameter F' . We found that the best choice of F' depended on the per-trace average disk access times. For those traces for which the average disk access time was small, in the 3-4ms range, it was best to take $F' = F$. For those traces for which the average disk access time was larger, it was best to take $F' = 4F$. This is not hard to explain. Traces with disk access times in the 3-4ms range must contain a great deal of sequential access, so that most requests hit in the disk's readahead cache and are served by the CSCAN scheduler in the order in which they are received. When this happens, it is not necessary to prefetch aggressively. When the disk access times are large, the access pattern is more complicated, and disk access times more varied. *Forestall's* mechanism for deciding when to prefetch benefits from overestimating the potential to stall. This smooths out the variations and avoids stalling due to the reordering of requests by CSCAN. Our implementation of *forestall* adapts to the observed disk access times, using the small value of F' for small disk access times (less than 5ms on average), and the larger value of F' for larger disk access times. Finally, because of the reordering of requests by CSCAN, we found it necessary to add *fixed horizon's* rule to issue a fetch whenever the cursor is within H requests of a missing block. This avoids stalling on reordered requests in situations in which the $iF' \geq d_i$ rule delays fetching until the cursor is very near the first missing block.

Rather than using complete lookahead information in our implementation of *forestall*, we check the value of the expression $iF - d_i$ only for those missing blocks within distance $2K$ of the cursor, where K is the cache size. We have not experimented with different values of this parameter, nor with variations of the history length 100 used to track fetch times and application process compute times.

Forestall's dependence on batch-size is similar to *aggressive's*. We used for *forestall* the batch sizes given in table 6.

To compare static vs. dynamic estimation of *forestall's* parameter F' , we compared the performance of *forestall* using fixed values of 1, 2, 4, 8, 15, 30, and 60 for F' to its performance using the dynamic estimation just described. Because actual average inter-reference compute times in our traces vary greatly (from 1.3ms for postgres-join to 15.7ms for the postgres-select), no single value can work well for all traces.




Figure 8: Performance on the synth (left) and xds (right) traces. Each group of bars represents the performance of the three algorithms *fixed horizon*, *aggressive*, and *forestall*, in left-to-right order.




Figure 9: Performance on the cscope2 trace.

The values with the least maximum degradation relative to the dynamic algorithm, over all traces and disk array sizes, are 30 and 60; performance is at most 6.8% worse than the dynamic estimation (on the j2 trace with 2-6 disks). We exclude the synth trace from this calculation, since its artificially low disk response times demand a very low value of F' . For each trace, there is a fixed choice of F that works well across all disk array sizes. This best choice varies from 1 for the dinero trace to 60 for the glimpse trace. The maximum degradation relative to the dynamic estimation allowing this much flexibility is 1.4% (for the ld trace with 7 disks and $F' = 30$). Finally, if we choose the best fixed value for each trace and each disk array size, the maximum degradation relative to the dynamic estimation is 1.2% for the cscope1 trace with one disk and $F' = 2$; for all other traces and array sizes, as well as all other array sizes for this trace, the degradation is less than .5%.

These results indicate that choosing the right parameters between workloads is more important than choosing the right parameter within a particular workload. Furthermore, *forestall*'s performance even with a single fixed parameter over all workloads and array sizes is always within 7% of optimal, and is almost always within 4% of optimal. This suggests that the advantages of *forestall* are due to its estimation of and adaptivity to upcoming disk load rather than the dynamic nature of its fetch-time and compute-time estimates.

5.1 Performance of *forestall*

Figure 8 shows the performance of the three practical algorithms, *fixed horizon*, *aggressive* and *forestall*, on the synthetic trace and xds. *Forestall* behaves exactly as expected. In the I/O bound situations, it prefetches aggressively enough to perform as well as or even better than *aggressive*. In the CPU-bound situations, it becomes more conservative in its prefetching, and has a lower driver overhead, matching the performance of *fixed horizon*.

Figures 9 and 10 show the performance of the three algorithms on the cscope2 and glimpse traces. Once again, *forestall* has the best performance of the three practical algorithms. On all remaining traces, over all configurations, *forestall*'s performance was between 2% worse and 5.8% better than the best of *aggressive* and *fixed horizon* in that configuration. For the full data, see the appendix.

Table 8 shows the utilization of the disks by *forestall* on the postgres-select trace. Its utilization falls between those of *aggressive* and *fixed horizon*, as expected. Moreover, in I/O-bound situations, it places a load on the disks similar to *aggressive*'s; in compute-bound situations, it places a lower load on the disks, similar to that of *fixed horizon*.




Figure 10: Performance on the glimpse trace.

disks	1	2	3	4	5	6
util.	.99	.92	.87	.81	.68	.63
disks	7	8	10	12	16	
util.	.62	.54	.39	.30	.32	

Table 8: Utilization of disks by *forestall* on the postgres-select trace.

6 Conclusions

This paper presents the results of a trace-driven simulation study of integrated prefetching and caching algorithms on a single read-only access sequence, assuming that all accesses are known in advance. We studied four algorithms: *aggressive*, *fixed horizon*, *reverse aggressive*, and *forestall*. We found that the theoretically near-optimal *reverse aggressive* usually has the best performance of the four algorithms, but that, perhaps surprisingly, it was never much better than the best of the other algorithms. This shows that carefully choosing replacements is not necessary to balance the load across the disks when the data is well laid out. We found that each of *aggressive* and *fixed horizon* performs well under the conditions for which it was designed, and in any given situation, one or the other performs similarly to *reverse aggressive*. Clearly, *aggressive* and *fixed horizon* are much more practical algorithms than *reverse aggressive*. These observations led us to the hybrid approach of *forestall*, which prefetches more aggressively in I/O-bound situations and more conservatively in compute-bound situations, resulting in nearly the best performance of the four in all configurations.

This study leaves several important issues unresolved. The performance of the algorithms depends on a set of parameters which interact in a complicated way with the applications’ access patterns and inter-reference compute times, the layout of data on disks, the disk-scheduling discipline, and the characteristics of the disks. At this time, we have no analytical basis for dynamically determining *aggressive*’s batch size, *fixed horizon*’s prefetch horizon H , *reverse aggressive*’s batch sizes and estimate of F , or *forestall*’s batch size and estimate F' of F . It is a challenging open problem to fully understand the interaction between the algorithmic parameters and the specific application and system charac-

teristics.

Another direction for future research is the treatment of writes, both theoretically and experimentally.

We have not considered the effects of incomplete or inaccurate hints and we have not dealt with the question of how to allocate buffers among competing processes. While the three practical prefetching algorithms can easily be adapted to deal with these situations ([5, 25]), we expect differences in their performance. Aggressive prefetching increases both disk utilization and cache utilization. Therefore, disks are more likely to be busy when unhinted accesses occur. Moreover, an aggressively prefetching process might consume too large a fraction of the cache relative to a nonhinting process. Since *fixed horizon* places the least load on the disks and the cache, it is likely to be least affected by unhinted accesses and to have the smallest impact on other executing processes.

Lastly, this work reaffirms that the operating system can effectively take advantage of hints. An important research direction is to determine how applications can easily provide such hints.

Acknowledgements

Tracy Kimbrel and Anna Karlin wish to thank Martin Tompa for his continued encouragement and advice. We could not have managed the production of this document without the help of Dylan McNamee. Karin Petersen, our paper shepherd, was of great help in improving the quality of presentation.

This research is supported in part by NSF grant numbers ECD-8907068, CCR-9301186, GER-9450075, CCR-9632769, in part by DARPA Contract numbers DABT63-94-C-0049, DABT63-93-C-0054, in part by generous contributions from the member companies of the Parallel Data Consortium, and in part by Intel Corporation and Digital Equipment Corporation. Tracy Kimbrel is supported by an Intel Foundation Graduate Fellowship. Brian Bershad is supported by an NSF Presidential Faculty Fellowship. Ed Felten is supported by an NSF National Young Investigator Award. The views and conclusions contained in this document are those of the au-

thors and should not be interpreted as representing the official policies, either expressed or implied, of any supporting organization or the U.S. Government.

References

- [1] A. Acharya, M. Uysal, R. Bennett, A. Mendelson, M. Beynon, J. Hollingsworth, J. Saltz and A. Sussman. Tuning the Performance of I/O-Intensive Parallel Applications. *Proceedings of the Fourth Annual Workshop on I/O in Parallel and Distributed Systems*, pages 15–27, May, 1996.
- [2] L.A. Belady. A Study of Replacement Algorithms for Virtual Storage Computers. *IBM Systems Journal*, 5(2):78–101, 1966.
- [3] Jim Gray. *The Benchmark Handbook*. Morgan-Kaufman, San Mateo, CA. 1991.
- [4] Pei Cao, Edward Felten, and Kai Li. Application-Controlled File Caching Policies. In *USENIX Summer 1994 Technical Conference*, pages 171–182, June 1994.
- [5] Pei Cao, Edward W. Felten, Anna Karlin, and Kai Li. Implementation and Performance of Integrated Application-Controlled Caching, Prefetching and Disk Scheduling. Technical Report TR-CS95-493, Princeton University, 1995.
- [6] Pei Cao, Edward W. Felten, Anna Karlin, and Kai Li. A study of Integrated Prefetching and Caching Strategies. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pages 188–197, May 1995.
- [7] Pei Cao, Edward W. Felten, and Kai Li. Implementation and Performance of Application-Controlled File Caching. In *Proceedings of the First USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 165–178, November 1994.
- [8] P.M. Chen and D.A. Patterson. Maximizing Performance in a Striped Disk Array. In *Proceedings of the 17th Annual Symposium on Computer Architecture*, pages 322–331, May 1990.
- [9] H.T. Chou and D.J. DeWitt. An Evaluation of Buffer Management Strategies for Relational Database Systems. In *Proceedings of the 19th International Conference on Very Large Data Bases*, pages 127–141, Dublin, Ireland, 1993.
- [10] Kenneth M. Curewitz, P. Krishnan, and Jeffrey S. Vitter. Practical Prefetching via Data Compression. In *Proceedings of the 1993 ACM Conference on Management of Data (SIGMOD)*, pages 257–266, Washington, DC, May 1993.
- [11] Carla Schlatter Ellis and David Kotz. Prefetching in File System for MIMD Multiprocessors. In *Proceedings of the 1989 International Conference on Parallel Processing*, pages 306–314, August 1989.
- [12] R.J. Feiertag and E.I. Organisk. The Multics Input/Output System. In *Proceedings of the 3rd Symposium on Operating Systems Principles*, pages 35–41, 1971.
- [13] Jim Griffioen and Randy Appleton. Reducing File System Latency using a Predictive Approach. In *USENIX Summer 1994 Technical Conference*, pages 197–208, June 1994.
- [14] John H. Howard, Michael Kazar, Sherri G. Menees, David A. Nichols, M. Satyanarayanan, Robert N. Sidebotham, and Michael J. West. Scale and Performance in a Distributed File System. *ACM Transactions on Computer Systems*, 6(1):51–81, February 1988.
- [15] M. Kim. Synchronized Disk Interleaving. *IEEE Transactions on Computers*, 35(11):978–988, 1986.
- [16] Tracy Kimbrel and Anna R. Karlin. Near-optimal Parallel Prefetching and Caching. In *Proceedings of the 1996 IEEE Symposium on Foundations of Computer Science*, October 1996.
- [17] David Kotz and Carla Schlatter Ellis. Practical Prefetching Techniques for Multiprocessor File Systems. *Journal of Distributed and Parallel Databases*, 1(1):33–51, January 1993.
- [18] David Kotz, Song Bac Toh, and Sriram Radhakrishnan. A Detailed Simulation Model of the HP 97560 Disk Drive. Technical Report PCS-TR94-220, Department of Computer Science, Dartmouth College, July 1994.
- [19] Marshall K. McKusick, William N. Joy, Samuel J. Leffler, and Robert S. Fabry. A Fast File System for UNIX. *ACM Transactions on Computer Systems*, 2(3):181–197, August 1984.
- [20] L. W. McVoy and S. R. Kleiman. Extent-like Performance from a UNIX File System. In *Proceedings of the 1991 Winter USENIX Conference*, pages 33–43, 1991.
- [21] Michael N. Nelson, Brent B. Welch, and John K. Ousterhout. Caching in the Sprite File System. *ACM Transactions on Computer Systems*, 6(1):134–154, February 1988.
- [22] Mark Palmer and Stanley B. Zdonik. Fido: A Cache That Learns to Fetch. In *Proceedings of the 17th International Conference on Very Large Data Bases*, pages 255–264, September 1991.
- [23] D.A. Patterson, G. Gibson, and R.H. Katz. A Case for Redundant Arrays for Inexpensive Disks (RAID). In *Proceedings of ACM SIGMOD Conference*, pages 109–116, June 1988.
- [24] R. Hugo Patterson and Garth A. Gibson. Exposing I/O Concurrency with Informed Prefetching. In *Proceedings of the Third International Conference on Parallel and Distributed Information Systems*, pages 7–16, September 1994.
- [25] R.H. Patterson, G.A. Gibson, E. Ginting, D. Stodolsky, and J. Zelenka. Informed Prefetching and Caching. In *Proceedings of the 15th Symposium on Operating Systems Principles*, pages 79–95, December 1995.
- [26] Chris Ruemmler and John Wilkes. An Introduction to Disk Drive Modelling. In *IEEE Computer*, 27(3):17–28, March 1994.
- [27] K. Salem and H. Garcia-Molina. Disk Striping. In *the 2nd IEEE Conference on Data Engineering*, pages 336–342, Feb. 1986.
- [28] Alan J. Smith. Second Bibliography on Cache Memories. *Computer Architecture News*, 19(4):154–182, June 1991.
- [29] C. Tait and D. Duchamp. Service Interface and Replica Management Algorithm for Mobile File System Clients. In *Proceedings of Parallel and Distributed Information Systems*, pages 190–196. IEEE, 1991.

A Performance data: baseline measurements

This section contains the raw simulation data for the baseline parameters as described in section 4: the prefetch horizon of *fixed horizon* is 62, *aggressive*'s batch size is set according to table 6, *reverse aggressive*'s fetch time estimate F and batch size are chosen to minimize its elapsed time, and Forestall's fetch time estimate F' is determined dynamically as described in section 5.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	4771	4771	4771	4771	4771	4771
driver time (sec)	2.3855	2.3855	2.3855	2.3855	2.3855	2.3855
stall time (sec)	0.027	0.009	0.009	0.009	0.009	0.009
elapsed time (sec)	105.951	105.933	105.933	105.933	105.933	105.933
average fetch time (msec)	3.156	3.178	3.233	3.258	3.274	3.319
average disk utilization	0.14	0.072	0.049	0.037	0.029	0.025
Aggressive						
fetches	8812	8812	8823	8815	8812	8816
driver time (sec)	4.406	4.406	4.4115	4.4075	4.406	4.408
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	107.944	107.95	107.946	107.944	107.947
average fetch time (msec)	3.141	3.146	3.174	3.176	3.188	3.203
average disk utilization	0.26	0.13	0.086	0.065	0.052	0.044
Reverse Aggressive						
fetches	4731	4764	4829	4830	4914	5018
driver time (sec)	2.3655	2.382	2.4145	2.415	2.457	2.509
stall time (sec)	0.023	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	105.927	105.941	105.972	105.97	106.01	106.06
average fetch time (msec)	3.31	3.36	3.366	3.437	3.494	3.288
average disk utilization	0.15	0.076	0.051	0.039	0.032	0.026
Forestall						
fetches	4753	4753	4753	4753	4753	4753
driver time (sec)	2.3765	2.3765	2.3765	2.3765	2.3765	2.3765
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	106.06	105.915	105.915	105.915	105.915	105.916
average fetch time (msec)	3.183	3.196	3.253	3.272	3.298	3.324
average disk utilization	0.14	0.072	0.049	0.037	0.03	0.025

Table 9: Performance on the dinero trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	4953	4953	4953	4953	4953	4953
driver time (sec)	2.4765	2.4765	2.4765	2.4765	2.4765	2.4765
stall time (sec)	3.131	0.013	0.013	0.013	0.013	0.013
elapsed time (sec)	30.542	27.424	27.424	27.424	27.424	27.424
average fetch time (msec)	3.53	3.239	3.248	3.286	3.317	3.355
average disk utilization	0.57	0.29	0.2	0.15	0.12	0.1
Aggressive						
fetches	6931	8570	8672	8678	8621	8576
driver time (sec)	3.4655	4.285	4.336	4.339	4.3105	4.288
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.219	29.27	29.273	29.245	29.223
average fetch time (msec)	3.758	3.361	3.429	3.365	3.39	3.356
average disk utilization	0.89	0.49	0.34	0.25	0.2	0.16
Reverse Aggressive						
fetches	5349	4995	5024	5093	5132	5135
driver time (sec)	2.6745	2.4975	2.512	2.5465	2.566	2.5675
stall time (sec)	1.312	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	28.921	27.453	27.465	27.498	27.515	27.515
average fetch time (msec)	3.622	3.344	3.376	3.409	3.396	3.618
average disk utilization	0.67	0.3	0.21	0.16	0.13	0.11
Forestall						
fetches	5210	4970	4953	4953	4953	4953
driver time (sec)	2.605	2.485	2.4765	2.4765	2.4765	2.4765
stall time (sec)	1.266	0	0	0	0	0.001
elapsed time (sec)	28.805	27.419	27.411	27.411	27.411	27.412
average fetch time (msec)	3.794	3.334	3.276	3.295	3.326	3.342
average disk utilization	0.69	0.3	0.2	0.15	0.12	0.1

Table 10: Performance on the cscope1 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	32.802	22.261	14.616	9.04	5.921	3.905	2.488	1.347	1.016	0.371	0.133
elapsed time (sec)	72.894	62.353	54.708	49.132	46.013	43.997	42.58	41.439	41.108	40.463	40.225
average fetch time (msec)	9.469	15.009	17.309	17.993	18.463	18.921	18.894	19.083	19.216	19.217	19.542
average disk utilization	0.77	0.72	0.63	0.55	0.48	0.43	0.38	0.34	0.28	0.24	0.18
Aggressive											
fetches	6318	6592	8208	8956	10299	11014	11587	11717	11619	11102	10662
driver time (sec)	3.159	3.296	4.104	4.478	5.1495	5.507	5.7935	5.8585	5.8095	5.551	5.331
stall time (sec)	15.858	5.597	1.798	0	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	56.126	46.002	43.011	41.587	42.259	42.617	42.903	42.977	42.924	42.661	42.44
average fetch time (msec)	8.773	13.256	14.354	16.514	17.138	17.683	17.722	17.551	17.224	17.65	18.201
average disk utilization	0.99	0.95	0.91	0.89	0.84	0.76	0.68	0.6	0.47	0.38	0.29
Reverse Aggressive											
fetches	6359	7320	6837	6290	6124	6071	6085	6115	6131	6177	6237
driver time (sec)	3.1795	3.66	3.4185	3.145	3.062	3.0355	3.0425	3.0575	3.0655	3.0885	3.1185
stall time (sec)	17.966	6.057	0.978	0	0.005	0.013	0.011	0.009	0.005	0.016	0.008
elapsed time (sec)	58.255	46.826	41.506	40.254	40.176	40.158	40.163	40.176	40.18	40.214	40.236
average fetch time (msec)	8.173	11.43	13.428	16.847	17.651	17.939	18.64	18.616	19.054	19.133	19.285
average disk utilization	0.89	0.89	0.74	0.66	0.54	0.45	0.4	0.35	0.29	0.24	0.19
Forestall											
fetches	6318	6467	7217	7239	7715	7387	7355	7187	7086	6853	6476
driver time (sec)	3.159	3.2335	3.6085	3.6195	3.8575	3.6935	3.6775	3.5935	3.543	3.4265	3.238
stall time (sec)	15.858	5.677	1.798	0	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	56.126	46.02	42.516	40.729	40.967	40.804	40.787	40.712	40.657	40.537	40.347
average fetch time (msec)	8.773	13.251	14.466	16.675	16.97	18.158	18.274	18.821	19.133	19.123	19.23
average disk utilization	0.99	0.93	0.82	0.74	0.64	0.55	0.47	0.42	0.33	0.27	0.19

Table 11: Performance on the cscope2 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739
driver time (sec)	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	28.459	12.906	7.046	2.961	1.669	0.762	0.221	0.164	0.152	0.014	0.014
elapsed time (sec)	108.429	92.876	87.016	82.931	81.639	80.732	80.191	80.134	80.122	79.984	79.984
average fetch time (msec)	7.843	11.914	14.814	16.147	16.993	17.482	17.906	18.178	18.671	18.875	19.108
average disk utilization	0.85	0.75	0.67	0.57	0.49	0.42	0.37	0.33	0.27	0.23	0.18
Aggressive											
fetches	12092	13572	15938	16740	17713	18081	17894	17577	16917	16542	16314
driver time (sec)	6.046	6.786	7.969	8.37	8.8565	9.0405	8.947	8.7885	8.4585	8.271	8.157
stall time (sec)	13.943	2.862	0.64	0.052	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	94.09	83.749	82.71	82.523	82.957	83.142	83.048	82.898	82.564	82.373	82.258
average fetch time (msec)	7.741	11.597	14.215	15.92	16.553	16.568	16.711	16.905	17.605	17.966	18.49
average disk utilization	0.99	0.94	0.91	0.81	0.71	0.6	0.51	0.45	0.36	0.3	0.23
Reverse Aggressive											
fetches	12228	12814	12501	12033	11880	11837	11852	11883	11919	11954	12004
driver time (sec)	6.114	6.407	6.2505	6.0165	5.94	5.9185	5.926	5.9415	5.9595	5.977	6.002
stall time (sec)	23.85	3.531	0.66	0.407	0.006	0.013	0.011	0.009	0.005	0.016	0.008
elapsed time (sec)	104.065	84.039	81.011	80.524	80.047	80.032	80.038	80.051	80.065	80.094	80.111
average fetch time (msec)	7.763	10.787	14.095	15.97	16.612	17.358	17.844	18.127	18.482	18.749	19.025
average disk utilization	0.91	0.82	0.73	0.6	0.49	0.43	0.38	0.34	0.28	0.23	0.18
Forestall											
fetches	12054	13115	14217	13969	14125	13878	13846	13589	13322	13048	12536
driver time (sec)	6.027	6.5575	7.1085	6.9845	7.0625	6.939	6.923	6.7945	6.661	6.524	6.268
stall time (sec)	14.273	2.863	0.64	0.052	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	94.401	83.521	81.849	81.137	81.163	81.041	81.024	80.904	80.767	80.626	80.369
average fetch time (msec)	7.731	11.603	13.616	15.745	16.183	17.246	17.648	18.477	18.669	18.812	19.035
average disk utilization	0.99	0.91	0.79	0.68	0.56	0.49	0.43	0.39	0.31	0.25	0.19

Table 12: Performance on the cscope3 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493
driver time (sec)	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465
stall time (sec)	65.619	31.046	20.054	14.029	10.381	7.886	5.702	4.769	2.809	1.404	0.722
elapsed time (sec)	107.582	73.009	62.017	55.992	52.344	49.849	47.665	46.732	44.772	43.367	42.685
average fetch time (msec)	13.424	15.145	16.192	17.244	18.068	18.33	18.452	18.642	18.555	18.571	18.743
average disk utilization	0.81	0.67	0.57	0.5	0.45	0.4	0.36	0.32	0.27	0.23	0.18
Aggressive											
fetches	6690	6888	7287	7551	8908	9376	10423	10992	12009	11530	11315
driver time (sec)	3.345	3.444	3.6435	3.7755	4.454	4.688	5.2115	5.496	6.0045	5.765	5.6575
stall time (sec)	54.58	18.58	6.384	2.495	0.826	0.035	0	0.009	0.005	0.001	0
elapsed time (sec)	96.641	60.74	48.744	44.987	43.996	43.439	43.928	44.221	44.726	44.482	44.374
average fetch time (msec)	12.889	14.259	14.645	16.247	15.973	16.836	16.79	16.896	16.137	15.917	16.198
average disk utilization	0.89	0.81	0.73	0.68	0.65	0.61	0.57	0.52	0.43	0.34	0.26
Reverse Aggressive											
fetches	6712	7179	7630	8141	7619	6803	6656	6709	6750	6822	6978
driver time (sec)	3.356	3.5895	3.815	4.0705	3.8095	3.4015	3.328	3.3545	3.375	3.411	3.489
stall time (sec)	52.011	15.928	4.971	0.495	0	0	0.011	0.009	0.005	0.006	0
elapsed time (sec)	94.083	58.234	47.502	43.282	42.526	42.118	42.055	42.08	42.096	42.133	42.205
average fetch time (msec)	12.745	13.46	13.793	13.877	14.73	17.016	18.321	18.37	18.541	18.514	18.406
average disk utilization	0.91	0.83	0.74	0.65	0.53	0.46	0.41	0.37	0.3	0.25	0.19
Forestall											
fetches	6610	6617	6945	6905	7033	6937	7113	7093	7125	7089	6941
driver time (sec)	3.305	3.3085	3.4725	3.4525	3.5165	3.4685	3.5565	3.5465	3.5625	3.5445	3.4705
stall time (sec)	54.886	18.833	6.58	2.906	1.397	0.099	0	0.009	0.005	0.001	0
elapsed time (sec)	96.907	60.858	48.769	45.075	43.63	42.284	42.273	42.272	42.284	42.262	42.187
average fetch time (msec)	13.001	14.378	14.797	16.382	16.651	17.627	17.557	18.099	18.015	18.055	18.273
average disk utilization	0.89	0.78	0.7	0.63	0.54	0.48	0.42	0.38	0.3	0.25	0.19

Table 13: Performance on the glimpse trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	2904	2904	2904	2904	2904	2904	2904	2904	2904	2904	2904
driver time (sec)	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452
stall time (sec)	15.281	7.297	4.696	3.043	2.086	1.565	1.212	1.041	0.599	0.416	0.269
elapsed time (sec)	24.898	16.914	14.313	12.66	11.703	11.182	10.829	10.658	10.216	10.033	9.886
average fetch time (msec)	8.368	10.94	13.299	15.031	16.214	16.93	17.502	17.657	18.467	18.945	19.2
average disk utilization	0.98	0.94	0.9	0.86	0.8	0.73	0.67	0.6	0.52	0.46	0.35
Aggressive											
fetches	2981	2982	3137	3102	3310	3505	3734	3779	4091	4285	4651
driver time (sec)	1.4905	1.491	1.5685	1.551	1.655	1.7525	1.867	1.8895	2.0455	2.1425	2.3255
stall time (sec)	15.245	6.329	3.433	2.052	0.579	0.265	0.023	0.009	0.005	0.001	0
elapsed time (sec)	24.9	15.985	13.166	11.768	10.399	10.182	10.055	10.063	10.215	10.308	10.49
average fetch time (msec)	8.248	10.583	12.037	14.199	14.932	15.958	16.444	17.175	17.62	18.017	18.261
average disk utilization	0.99	0.99	0.96	0.94	0.95	0.92	0.87	0.81	0.71	0.62	0.51
Reverse Aggressive											
fetches	3041	3079	3202	3312	3161	3037	3103	3000	2953	3004	3008
driver time (sec)	1.5205	1.5395	1.601	1.656	1.5805	1.5185	1.5515	1.5	1.4765	1.502	1.504
stall time (sec)	14.662	6.217	3.233	1.704	0.879	0.618	0.211	0.151	0.035	0.016	0.008
elapsed time (sec)	24.347	15.921	12.999	11.525	10.624	10.301	9.927	9.816	9.676	9.683	9.677
average fetch time (msec)	7.932	10.036	11.585	13.254	14.43	16.016	16.269	17.39	18.558	18.995	18.992
average disk utilization	0.99	0.97	0.95	0.95	0.86	0.79	0.73	0.66	0.57	0.49	0.37
Forestall											
fetches	2981	2982	3137	3102	3310	3505	3734	3799	3896	3799	3147
driver time (sec)	1.4905	1.491	1.5685	1.551	1.655	1.7525	1.867	1.8995	1.948	1.8995	1.5735
stall time (sec)	15.245	6.329	3.433	2.052	0.579	0.265	0.023	0.013	0.005	0.001	0
elapsed time (sec)	24.9	15.985	13.166	11.768	10.399	10.182	10.055	10.077	10.118	10.065	9.738
average fetch time (msec)	8.248	10.583	12.037	14.199	14.932	15.958	16.442	17.253	18.126	18.582	18.983
average disk utilization	0.99	0.99	0.96	0.94	0.95	0.92	0.87	0.81	0.7	0.58	0.38

Table 14: Performance on the ld trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	3856	3856	3856	3856	3856	3856
driver time (sec)	1.928	1.928	1.928	1.928	1.928	1.928
stall time (sec)	4.723	0.04	0.017	0.017	0.017	0.017
elapsed time (sec)	85.867	81.184	81.161	81.161	81.161	81.161
average fetch time (msec)	17.228	18.029	18.039	18.299	18.344	18.094
average disk utilization	0.77	0.43	0.29	0.22	0.17	0.14
Aggressive						
fetches	4698	5836	6225	6156	6047	5919
driver time (sec)	2.349	2.918	3.1125	3.078	3.0235	2.9595
stall time (sec)	3.994	0.152	0.258	0	0	0.001
elapsed time (sec)	85.559	82.286	82.586	82.294	82.239	82.176
average fetch time (msec)	15.032	16.576	15.929	16.578	16.706	17.102
average disk utilization	0.83	0.59	0.4	0.31	0.25	0.21
Reverse Aggressive						
fetches	3987	3853	3859	3873	3879	3892
driver time (sec)	1.9935	1.9265	1.9295	1.9365	1.9395	1.946
stall time (sec)	3.775	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	84.984	81.163	81.164	81.169	81.17	81.175
average fetch time (msec)	15.776	18.055	17.964	18.204	18.262	17.934
average disk utilization	0.74	0.43	0.28	0.22	0.17	0.14
Forestall						
fetches	4694	4207	3929	3857	3855	3856
driver time (sec)	2.347	2.1035	1.9645	1.9285	1.9275	1.928
stall time (sec)	3.994	0.153	0.258	0	0	0.001
elapsed time (sec)	85.557	81.472	81.438	81.144	81.143	81.145
average fetch time (msec)	15.034	15.022	15.525	17.291	17.45	17.692
average disk utilization	0.82	0.39	0.25	0.21	0.17	0.14

Table 15: Performance on the postgres-join trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	32.37	12.647	5.943	3.154	1.402	0.581	0.476	0.073	0.034	0.018	0.018
elapsed time (sec)	45.39	25.667	18.963	16.174	14.422	13.601	13.496	13.093	13.054	13.038	13.038
average fetch time (msec)	14.368	14.906	15.044	15.13	15.347	15.413	15.437	15.411	15.278	15.356	15.071
average disk utilization	0.98	0.9	0.82	0.72	0.66	0.58	0.5	0.45	0.36	0.3	0.22
Aggressive											
fetches	3085	3085	3085	3085	3286	3317	3826	3937	3902	3852	3731
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.643	1.6585	1.913	1.9685	1.951	1.926	1.8655
stall time (sec)	30.691	10.772	3.517	0.844	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.792	16.537	13.864	13.121	13.137	13.391	13.455	13.434	13.405	13.343
average fetch time (msec)	13.985	14.173	13.95	14.55	13.923	15.036	15.221	15.274	14.797	14.8	15.155
average disk utilization	0.99	0.92	0.87	0.81	0.7	0.63	0.62	0.56	0.43	0.35	0.26
Reverse Aggressive											
fetches	3106	3106	3318	3110	3109	3108	3112	3122	3116	3122	3124
driver time (sec)	1.553	1.553	1.659	1.555	1.5545	1.554	1.556	1.561	1.558	1.561	1.562
stall time (sec)	28.956	8.461	2.66	0.125	0	0.001	0	0	0	0	0.002
elapsed time (sec)	41.987	21.492	15.797	13.158	13.032	13.033	13.034	13.039	13.036	13.039	13.042
average fetch time (msec)	13.248	12.704	12.127	13.581	14.394	15.051	14.803	14.049	14.618	14.024	14.143
average disk utilization	0.98	0.92	0.85	0.8	0.69	0.6	0.5	0.42	0.35	0.28	0.21
Fore stall											
fetches	3085	3085	3085	3085	3085	3305	3797	3795	3399	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.6525	1.8985	1.8975	1.6995	1.5425	1.5425
stall time (sec)	30.691	10.791	3.517	0.844	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.811	16.537	13.864	13.02	13.131	13.376	13.384	13.182	13.021	13.02
average fetch time (msec)	13.985	14.154	13.933	14.524	14.392	15.056	15.242	15.249	15.2	15.086	15.032
average disk utilization	0.99	0.92	0.87	0.81	0.68	0.63	0.62	0.54	0.39	0.3	0.22

Table 16: Performance on the postgres-select trace.

Disks	1	2	3	4
Fixed Horizon				
fetches	38000	38000	38000	38000
driver time (sec)	19	19	19	19
stall time (sec)	82.583	12.044	0	0
elapsed time (sec)	201.439	130.9	118.856	118.856
average fetch time (msec)	3.748	3.776	3.229	3.214
average disk utilization	0.71	0.55	0.34	0.26
Aggressive				
fetches	39240	41902	100994	100548
driver time (sec)	19.62	20.951	50.497	50.274
stall time (sec)	36.37	0.933	0.015	0.015
elapsed time (sec)	155.846	121.74	150.368	150.145
average fetch time (msec)	3.965	5.647	3.37	3.164
average disk utilization	1	0.97	0.75	0.53
Reverse Aggressive				
fetches	39265	42000	37907	38148
driver time (sec)	19.6325	21	18.9535	19.074
stall time (sec)	41.599	2.765	0.014	0.015
elapsed time (sec)	161.088	123.621	118.824	118.945
average fetch time (msec)	3.928	3.907	3.762	3.958
average disk utilization	0.96	0.66	0.4	0.32
Fore stall				
fetches	39240	38900	39838	38000
driver time (sec)	19.62	19.45	19.919	19
stall time (sec)	36.37	1.232	0.016	0
elapsed time (sec)	155.846	120.538	119.791	118.856
average fetch time (msec)	3.965	4.895	4.843	3.218
average disk utilization	1	0.79	0.54	0.26

Table 17: Performance on the synth trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	5900	5900	5900	5900	5900	5900
driver time (sec)	2.95	2.95	2.95	2.95	2.95	2.95
stall time (sec)	32.582	4.964	3.219	1.138	0.474	0.094
elapsed time (sec)	65.611	37.993	36.248	34.167	33.503	33.123
average fetch time (msec)	10.74	7.758	14.065	10.106	15.07	10.869
average disk utilization	0.97	0.6	0.76	0.44	0.53	0.32
Aggressive						
fetches	5925	7778	6563	9831	8312	10215
driver time (sec)	2.9625	3.889	3.2815	4.9155	4.156	5.1075
stall time (sec)	30.667	0.337	0.356	0.129	0.133	0.055
elapsed time (sec)	63.708	34.305	33.716	35.123	34.368	35.241
average fetch time (msec)	10.711	7.496	14.101	9.801	15.454	10.711
average disk utilization	1	0.85	0.91	0.69	0.75	0.52
Reverse Aggressive						
fetches	5892	5989	5927	6001	5893	6017
driver time (sec)	2.946	2.9945	2.9635	3.0005	2.9465	3.0085
stall time (sec)	31.155	0.275	0.528	0.046	0.017	0.018
elapsed time (sec)	64.18	33.348	33.57	33.125	33.042	33.105
average fetch time (msec)	10.79	7.732	14.092	9.864	14.883	10.173
average disk utilization	0.99	0.69	0.83	0.45	0.53	0.31
Forestall						
fetches	5925	6929	6553	7451	7882	7032
driver time (sec)	2.9625	3.4645	3.2765	3.7255	3.941	3.516
stall time (sec)	30.667	0.337	0.356	0.129	0.133	0.055
elapsed time (sec)	63.708	33.88	33.711	33.933	34.153	33.65
average fetch time (msec)	10.711	7.559	14.082	9.945	15.53	10.7
average disk utilization	1	0.77	0.91	0.55	0.72	0.37

Table 18: Performance on the xds trace.

B Performance data: FCFS

This section contains the data for the baseline parameters as in the previous section, but with FCFS disk head scheduling rather than CSCAN.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	4771	4771	4771	4771	4771	4771
driver time (sec)	2.3855	2.3855	2.3855	2.3855	2.3855	2.3855
stall time (sec)	0.009	0.009	0.009	0.009	0.009	0.009
elapsed time (sec)	105.933	105.933	105.933	105.933	105.933	105.933
average fetch time (msec)	3.153	3.181	3.218	3.247	3.26	3.314
average disk utilization	0.14	0.072	0.048	0.037	0.029	0.025
Aggressive						
fetches	8812	8812	8812	8814	8812	8814
driver time (sec)	4.406	4.406	4.406	4.407	4.406	4.407
stall time (sec)	0	0	0	0	0	0.001
elapsed time (sec)	107.944	107.944	107.944	107.945	107.944	107.946
average fetch time (msec)	3.148	3.144	3.166	3.169	3.178	3.194
average disk utilization	0.26	0.13	0.086	0.065	0.052	0.043
Reverse Aggressive						
fetches	4731	4764	4829	4830	4914	5018
driver time (sec)	2.3655	2.382	2.4145	2.415	2.457	2.509
stall time (sec)	0.023	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	105.927	105.941	105.972	105.97	106.01	106.06
average fetch time (msec)	3.311	3.356	3.352	3.431	3.483	3.284
average disk utilization	0.15	0.075	0.051	0.039	0.032	0.026

Table 19: Performance on the dinero trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	4953	4953	4953	4953	4953	4953
driver time (sec)	2.4765	2.4765	2.4765	2.4765	2.4765	2.4765
stall time (sec)	3.131	0.013	0.013	0.013	0.013	0.013
elapsed time (sec)	30.542	27.424	27.424	27.424	27.424	27.424
average fetch time (msec)	3.533	3.245	3.247	3.277	3.305	3.335
average disk utilization	0.57	0.29	0.2	0.15	0.12	0.1
Aggressive						
fetches	6778	8582	8606	8685	8621	8576
driver time (sec)	3.389	4.291	4.303	4.3425	4.3105	4.288
stall time (sec)	0.609	0	0	0	0	0.001
elapsed time (sec)	28.932	29.225	29.237	29.277	29.245	29.223
average fetch time (msec)	3.808	3.405	3.373	3.358	3.377	3.341
average disk utilization	0.89	0.5	0.33	0.25	0.2	0.16
Reverse Aggressive						
fetches	5349	4995	5024	5093	5132	5135
driver time (sec)	2.6745	2.4975	2.512	2.5465	2.566	2.5675
stall time (sec)	1.162	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	28.771	27.453	27.465	27.498	27.515	27.515
average fetch time (msec)	3.656	3.342	3.366	3.4	3.382	3.619
average disk utilization	0.68	0.3	0.21	0.16	0.13	0.11

Table 20: Performance on the cscopel trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	35.261	24.715	15.262	7.728	4.287	2.408	1.273	0.741	0.555	0.087	0.03
elapsed time (sec)	75.353	64.807	55.354	47.82	44.379	42.5	41.365	40.833	40.647	40.179	40.122
average fetch time (msec)	9.887	15.936	17.771	18.352	18.473	18.93	18.913	19.165	19.128	19.29	19.404
average disk utilization	0.78	0.73	0.64	0.57	0.5	0.44	0.39	0.35	0.28	0.24	0.18
Aggressive											
fetches	6196	6324	7302	8450	9933	10777	11475	11707	11529	11102	10662
driver time (sec)	3.098	3.162	3.651	4.225	4.9665	5.3885	5.7375	5.8535	5.7645	5.551	5.331
stall time (sec)	17.951	8.281	3.024	0	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	58.158	48.552	43.784	41.334	42.076	42.499	42.847	42.972	42.879	42.661	42.44
average fetch time (msec)	9.358	14.89	16.761	17.499	17.804	17.917	17.767	17.571	17.348	17.734	18.257
average disk utilization	1	0.97	0.93	0.89	0.84	0.76	0.68	0.6	0.47	0.38	0.29
Reverse Aggressive											
fetches	6359	7320	6837	6290	6124	6071	6085	6115	6131	6177	6237
driver time (sec)	3.1795	3.66	3.4185	3.145	3.062	3.0355	3.0425	3.0575	3.0655	3.0885	3.1185
stall time (sec)	19.611	12.595	3.118	0	0.017	0.013	0.011	0.009	0.005	0.016	0.008
elapsed time (sec)	59.9	53.364	43.646	40.254	40.188	40.158	40.163	40.176	40.18	40.214	40.236
average fetch time (msec)	8.869	13.946	16.125	17.786	18.115	18.048	18.615	18.678	19.186	19.328	19.283
average disk utilization	0.94	0.96	0.84	0.69	0.55	0.45	0.4	0.36	0.29	0.25	0.19

Table 21: Performance on the cscope2 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739
driver time (sec)	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	31.757	13.743	6.463	2.014	0.749	0.233	0.2	0.014	0.095	0.014	0.014
elapsed time (sec)	111.727	93.713	86.433	81.984	80.719	80.203	80.17	79.984	80.065	79.984	79.984
average fetch time (msec)	8.184	12.268	15.013	16.108	17.015	17.456	17.891	18.277	18.654	18.88	19.1
average disk utilization	0.86	0.77	0.68	0.58	0.49	0.43	0.37	0.34	0.27	0.23	0.18
Aggressive											
fetches	11974	12937	15104	16457	17588	18048	17824	17547	16917	16542	16254
driver time (sec)	5.987	6.4685	7.552	8.2285	8.794	9.024	8.912	8.7735	8.4585	8.271	8.127
stall time (sec)	18.727	4.342	1.132	0.107	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	98.815	84.911	82.785	82.436	82.895	83.126	83.013	82.883	82.564	82.373	82.228
average fetch time (msec)	8.219	12.54	15.309	16.142	16.775	16.717	16.777	16.966	17.597	17.96	18.603
average disk utilization	1	0.96	0.93	0.81	0.71	0.6	0.51	0.45	0.36	0.3	0.23
Reverse Aggressive											
fetches	12228	12814	12501	12033	11880	11837	11852	11883	11919	11954	12004
driver time (sec)	6.114	6.407	6.2505	6.0165	5.94	5.9185	5.926	5.9415	5.9595	5.977	6.002
stall time (sec)	27.37	6.138	1.178	0.195	0.005	0.013	0.011	0.009	0.005	0.016	0.008
elapsed time (sec)	107.585	86.646	81.529	80.312	80.046	80.032	80.038	80.051	80.065	80.094	80.111
average fetch time (msec)	8.408	12.53	15.639	16.608	16.992	17.386	17.832	18.199	18.542	18.83	19.076
average disk utilization	0.96	0.93	0.8	0.62	0.5	0.43	0.38	0.34	0.28	0.23	0.18

Table 22: Performance on the cscope3 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493
driver time (sec)	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465
stall time (sec)	69.337	31.846	19.994	12.419	8.631	6.095	4.014	3.171	1.547	0.745	0.388
elapsed time (sec)	111.3	73.809	61.957	54.382	50.594	48.058	45.977	45.134	43.51	42.708	42.351
average fetch time (msec)	14.011	15.68	16.63	17.531	18.305	18.515	18.611	18.648	18.529	18.586	18.675
average disk utilization	0.82	0.69	0.58	0.52	0.47	0.42	0.38	0.34	0.28	0.24	0.18
Aggressive											
fetches	6690	6786	7277	7400	8361	9196	10233	10959	11953	11500	11253
driver time (sec)	3.345	3.393	3.6385	3.7	4.1805	4.598	5.1165	5.4795	5.9765	5.75	5.6265
stall time (sec)	59.899	23.204	10.316	3.068	1.321	0.252	0	0.009	0.005	0.001	0
elapsed time (sec)	101.96	65.313	52.671	45.484	44.218	43.566	43.833	44.205	44.698	44.467	44.343
average fetch time (msec)	13.814	15.675	16.321	17.255	17.756	17.404	17.228	17.12	16.022	16.038	16.341
average disk utilization	0.91	0.81	0.75	0.7	0.67	0.61	0.57	0.53	0.43	0.35	0.26
Reverse Aggressive											
fetches	6712	7179	7630	8141	7619	6803	6656	6709	6750	6822	6978
driver time (sec)	3.356	3.5895	3.815	4.0705	3.8095	3.4015	3.328	3.3545	3.375	3.411	3.489
stall time (sec)	59.04	22.734	8.948	1.776	1.094	0	0.011	0.009	0.005	0.006	0
elapsed time (sec)	101.112	65.04	51.479	44.563	43.62	42.118	42.055	42.08	42.096	42.133	42.205
average fetch time (msec)	14.018	15.73	16.272	16.806	18.093	18.293	18.44	18.529	18.633	18.563	18.547
average disk utilization	0.93	0.87	0.8	0.77	0.63	0.49	0.42	0.37	0.3	0.25	0.19

Table 23: Performance on the glimpse trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	2904	2904	2904	2904	2904	2904	2904	2904	2904	2904	2904
driver time (sec)	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452	1.452
stall time (sec)	16.029	7.491	4.39	2.843	1.865	1.291	0.913	0.839	0.431	0.218	0.176
elapsed time (sec)	25.646	17.108	14.007	12.46	11.482	10.908	10.53	10.456	10.048	9.835	9.793
average fetch time (msec)	8.63	11.143	13.432	15.137	16.323	17.033	17.311	17.876	18.615	18.958	19.23
average disk utilization	0.98	0.95	0.93	0.88	0.83	0.76	0.68	0.62	0.54	0.47	0.36
Aggressive											
fetches	2943	2979	3000	3052	3228	3459	3712	3809	4061	4285	4608
driver time (sec)	1.4715	1.4895	1.5	1.526	1.614	1.7295	1.856	1.9045	2.0305	2.1425	2.304
stall time (sec)	15.651	7.305	3.955	2.126	0.983	0.509	0.017	0.009	0.005	0.001	0
elapsed time (sec)	25.287	16.959	13.62	11.817	10.762	10.403	10.038	10.078	10.2	10.308	10.469
average fetch time (msec)	8.575	11.002	13.144	14.781	15.716	16.17	16.718	16.973	17.723	18.074	18.23
average disk utilization	1	0.97	0.97	0.95	0.94	0.9	0.88	0.8	0.71	0.63	0.5
Reverse Aggressive											
fetches	3041	3079	3202	3312	3161	3037	3103	3000	2953	3004	3008
driver time (sec)	1.5205	1.5395	1.601	1.656	1.5805	1.5185	1.5515	1.5	1.4765	1.502	1.504
stall time (sec)	15.977	7.36	4.121	2.409	1.197	0.82	0.412	0.2	0.035	0.016	0.008
elapsed time (sec)	25.662	17.064	13.887	12.23	10.942	10.503	10.128	9.865	9.676	9.683	9.677
average fetch time (msec)	8.425	10.763	12.743	14.394	15.897	16.876	17.324	17.799	18.616	19.074	19.137
average disk utilization	1	0.97	0.98	0.97	0.92	0.81	0.76	0.68	0.57	0.49	0.37

Table 24: Performance on the ld trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	3856	3856	3856	3856	3856	3856
driver time (sec)	1.928	1.928	1.928	1.928	1.928	1.928
stall time (sec)	8.916	0.017	0.017	0.017	0.017	0.017
elapsed time (sec)	90.06	81.161	81.161	81.161	81.161	81.161
average fetch time (msec)	18.516	18.188	18.122	18.342	18.369	18.109
average disk utilization	0.79	0.43	0.29	0.22	0.17	0.14
Aggressive						
fetches	4138	5704	6188	6156	5978	5949
driver time (sec)	2.069	2.852	3.094	3.078	2.989	2.9745
stall time (sec)	8.546	0	0	0	0	0.001
elapsed time (sec)	89.831	82.068	82.31	82.294	82.205	82.191
average fetch time (msec)	18.584	17.688	16.933	17.105	17.293	17.357
average disk utilization	0.86	0.61	0.42	0.32	0.25	0.21
Reverse Aggressive						
fetches	3987	3853	3859	3873	3879	3892
driver time (sec)	1.9935	1.9265	1.9295	1.9365	1.9395	1.946
stall time (sec)	8.427	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	89.636	81.163	81.164	81.169	81.17	81.175
average fetch time (msec)	18.51	18.173	18.06	18.344	18.335	18.037
average disk utilization	0.82	0.43	0.29	0.22	0.18	0.14

Table 25: Performance on the postgres-join trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	39.131	13.893	6.434	3.239	1.313	0.489	0.186	0.018	0.028	0.018	0.018
elapsed time (sec)	52.151	26.913	19.454	16.259	14.333	13.509	13.206	13.038	13.048	13.038	13.038
average fetch time (msec)	16.703	16.165	16.017	15.85	15.708	15.656	15.632	15.526	15.495	15.414	15.117
average disk utilization	0.99	0.93	0.85	0.75	0.68	0.6	0.52	0.46	0.37	0.3	0.22
Aggressive											
fetches	3085	3085	3085	3085	3085	3234	3707	3891	3926	3882	3731
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.617	1.8535	1.9455	1.963	1.941	1.8655
stall time (sec)	39.072	13.453	4.9	1.342	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	52.092	26.473	17.92	14.362	13.02	13.096	13.331	13.432	13.446	13.42	13.343
average fetch time (msec)	16.703	16.214	16.06	15.91	15.742	15.512	15.774	15.251	14.928	14.962	15.202
average disk utilization	0.99	0.94	0.92	0.85	0.75	0.64	0.63	0.55	0.44	0.36	0.27
Reverse Aggressive											
fetches	3106	3106	3318	3110	3109	3108	3112	3122	3116	3122	3124
driver time (sec)	1.553	1.553	1.659	1.555	1.5545	1.554	1.556	1.561	1.558	1.561	1.562
stall time (sec)	39.031	13.202	5.806	1.008	0	0.001	0	0	0	0	0.002
elapsed time (sec)	52.062	26.233	18.943	14.041	13.032	13.033	13.034	13.039	13.036	13.039	13.042
average fetch time (msec)	16.628	16.156	16.057	15.859	15.694	15.667	15.619	15.497	15.469	15.426	15.262
average disk utilization	0.99	0.96	0.94	0.88	0.75	0.62	0.53	0.46	0.37	0.31	0.23

Table 26: Performance on the postgres-select trace.

Disks	1	2	3	4
Fixed Horizon				
fetches	38000	38000	38000	38000
driver time (sec)	19	19	19	19
stall time (sec)	82.583	12.044	0	0
elapsed time (sec)	201.439	130.9	118.856	118.856
average fetch time (msec)	3.748	3.776	3.229	3.214
average disk utilization	0.71	0.55	0.34	0.26
Aggressive				
fetches	39240	41902	100994	100548
driver time (sec)	19.62	20.951	50.497	50.274
stall time (sec)	36.37	0.933	0.015	0.015
elapsed time (sec)	155.846	121.74	150.368	150.145
average fetch time (msec)	3.965	5.647	3.37	3.164
average disk utilization	1	0.97	0.75	0.53
Reverse Aggressive				
fetches	39265	42000	37907	38148
driver time (sec)	19.6325	21	18.9535	19.074
stall time (sec)	41.599	2.765	0.014	0.015
elapsed time (sec)	161.088	123.621	118.824	118.945
average fetch time (msec)	3.928	3.907	3.762	3.958
average disk utilization	0.96	0.66	0.4	0.32

Table 27: Performance on the synth trace.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	5883	5883	5883	5883	5883	5883
driver time (sec)	2.9415	2.9415	2.9415	2.9415	2.9415	2.9415
stall time (sec)	34.937	8.068	4.974	2.068	1.005	0.242
elapsed time (sec)	68.644	41.775	38.681	35.775	34.712	33.949
average fetch time (msec)	10.86	7.769	14.127	10.142	14.97	10.835
average disk utilization	0.93	0.55	0.72	0.42	0.51	0.31
Aggressive						
fetches	5925	7662	6439	9847	8206	10114
driver time (sec)	2.9625	3.831	3.2195	4.9235	4.103	5.057
stall time (sec)	31.431	0.012	0	0	0	0
elapsed time (sec)	64.472	33.922	33.298	35.002	34.182	35.136
average fetch time (msec)	10.846	7.59	14.215	9.807	15.443	10.636
average disk utilization	1	0.86	0.92	0.69	0.74	0.51
Reverse Aggressive						
fetches	5910	5997	5945	6007	5904	6024
driver time (sec)	2.955	2.9985	2.9725	3.0035	2.952	3.012
stall time (sec)	31.298	0.043	0.334	0.01	0.012	0.011
elapsed time (sec)	65.018	33.807	34.072	33.779	33.729	33.788
average fetch time (msec)	10.913	7.729	14.213	9.877	15.115	10.124
average disk utilization	0.99	0.69	0.83	0.44	0.53	0.3

Table 28: Performance on the xds trace.

C Performance data: double-speed CPU

This section contains the data for the xds trace, with the processor speed doubled.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fixed Horizon											
fetches	5900	5900	5900	5900	5900	5900	5883	5883	5883	5883	5883
driver time (sec)	1.475	1.475	1.475	1.475	1.475	1.475	1.47075	1.47075	1.47075	1.47075	1.47075
stall time (sec)	47.186	16.135	13.901	7.272	6.058	2.778	5.577	5.343	3.038	3.306	2.317
elapsed time (sec)	63.698	32.647	30.413	23.784	22.57	19.29	22.422	22.188	19.883	20.151	19.162
average fetch time (msec)	10.714	7.735	14.014	10.122	15.41	10.959	15.97	12.872	12.847	14.048	13.583
average disk utilization	0.99	0.7	0.91	0.63	0.81	0.56	0.6	0.43	0.38	0.34	0.26
Aggressive											
fetches	5890	6272	5965	6471	5963	7602	6567	8701	10278	10607	10948
driver time (sec)	1.4725	1.568	1.49125	1.61775	1.49075	1.9005	1.64175	2.17525	2.5695	2.65175	2.737
stall time (sec)	46.755	10.572	11.808	2.384	2.776	0.048	0.097	0.045	0.046	0.036	0.026
elapsed time (sec)	63.264	27.177	28.336	19.038	19.303	16.985	17.113	17.595	17.99	18.062	18.137
average fetch time (msec)	10.737	7.664	14.121	10.064	15.496	10.871	16.398	12.921	12.7	13.709	12.392
average disk utilization	1	0.88	0.99	0.86	0.96	0.81	0.9	0.8	0.73	0.67	0.47
Reverse Aggressive											
fetches	5892	6095	5939	6182	6001	6017	5970	6042	6055	6090	6164
driver time (sec)	1.473	1.52375	1.48475	1.5455	1.50025	1.50425	1.4925	1.5105	1.51375	1.5225	1.541
stall time (sec)	47.078	8.592	11.946	1.135	2.836	0.055	0.304	0.045	0.046	0.036	0.034
elapsed time (sec)	63.588	25.152	28.467	17.717	19.373	16.596	17.171	16.93	16.934	16.933	16.949
average fetch time (msec)	10.787	7.734	14.132	10.228	15.506	10.787	16.032	12.257	11.954	12.902	12.794
average disk utilization	1	0.94	0.98	0.89	0.96	0.65	0.8	0.55	0.43	0.39	0.29

Table 29: Performance on the xds trace.

D Performance data: varying cache size

This section contains the data for several traces with cache sizes of 5MB (640 blocks) and 15 MB (1920 blocks).

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	7804	7804	7804	7804	7804	7804
driver time (sec)	3.902	3.902	3.902	3.902	3.902	3.902
stall time (sec)	80.295	38.262	24.527	16.624	12.193	9.234
elapsed time (sec)	122.913	80.88	67.145	59.242	54.811	51.852
average fetch time (msec)	13.65	15.451	16.454	17.39	18.142	18.402
average disk utilization	0.87	0.75	0.64	0.57	0.52	0.46
Aggressive						
fetches	8133	8407	9377	8777	9847	10220
driver time (sec)	4.0665	4.2035	4.6885	4.3885	4.9235	5.11
stall time (sec)	73.144	27.576	12.83	4.352	0.873	0.007
elapsed time (sec)	115.927	70.496	56.235	47.457	44.513	43.833
average fetch time (msec)	13.02	14.281	14.479	16.104	15.938	16.928
average disk utilization	0.91	0.85	0.8	0.74	0.71	0.66
Reverse Aggressive						
fetches	8280	9007	9967	9318	9802	8251
driver time (sec)	4.14	4.5035	4.9835	4.659	4.901	4.1255
stall time (sec)	69.892	24.098	7.214	1.796	0.064	0
elapsed time (sec)	112.748	67.318	50.914	45.171	43.681	42.842
average fetch time (msec)	12.649	13.017	13.017	13.804	14.166	17.142
average disk utilization	0.93	0.87	0.85	0.71	0.64	0.55

Table 30: Performance on the glimpse trace, cache size 640.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	5853	5853	5853	5853	5853	5853
driver time (sec)	2.9265	2.9265	2.9265	2.9265	2.9265	2.9265
stall time (sec)	58.697	27.363	18.29	12.82	9.592	7.442
elapsed time (sec)	100.34	69.006	59.933	54.463	51.235	49.085
average fetch time (msec)	13.302	14.879	16.005	17.096	17.981	18.251
average disk utilization	0.78	0.63	0.52	0.46	0.41	0.36
Aggressive						
fetches	6041	6121	6647	7044	8048	8390
driver time (sec)	3.0205	3.0605	3.3235	3.522	4.024	4.195
stall time (sec)	46.429	13.448	4.781	2.51	0.829	0.029
elapsed time (sec)	88.166	55.225	46.821	44.748	43.569	42.94
average fetch time (msec)	12.832	14.085	14.453	16.047	16.247	17.082
average disk utilization	0.88	0.78	0.68	0.63	0.6	0.56
Reverse Aggressive						
fetches	5998	6072	6542	6377	6441	6035
driver time (sec)	2.999	3.036	3.271	3.1885	3.2205	3.0175
stall time (sec)	42.301	11.025	3.324	0.964	0	0
elapsed time (sec)	84.016	52.777	45.311	42.869	41.937	41.734
average fetch time (msec)	12.749	13.675	13.685	14.398	14.649	17.02
average disk utilization	0.91	0.79	0.66	0.54	0.45	0.41

Table 31: Performance on the glimpse trace, cache size 1920.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	4640	4640	4640	4640	4640	4640
driver time (sec)	2.32	2.32	2.32	2.32	2.32	2.32
stall time (sec)	4.96	0.04	0.017	0.017	0.017	0.017
elapsed time (sec)	86.496	81.576	81.553	81.553	81.553	81.553
average fetch time (msec)	17.211	18.089	18.152	18.378	18.463	18.217
average disk utilization	0.92	0.51	0.34	0.26	0.21	0.17
Aggressive						
fetches	5378	7310	8013	8043	7758	7561
driver time (sec)	2.689	3.655	4.0065	4.0215	3.879	3.7805
stall time (sec)	3.994	0.152	0.258	0	0	0.001
elapsed time (sec)	85.899	83.023	83.48	83.237	83.095	82.997
average fetch time (msec)	15.152	17.332	16.677	16.709	16.899	16.909
average disk utilization	0.95	0.76	0.53	0.4	0.32	0.26
Reverse Aggressive						
fetches	4912	4615	4631	4657	4676	4691
driver time (sec)	2.456	2.3075	2.3155	2.3285	2.338	2.3455
stall time (sec)	3.655	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	85.327	81.544	81.55	81.561	81.569	81.574
average fetch time (msec)	16.064	18.147	18.068	18.361	18.282	18.026
average disk utilization	0.92	0.51	0.34	0.26	0.21	0.17

Table 32: Performance on the postgres-join trace, cache size 640.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	3793	3793	3793	3793	3793	3793
driver time (sec)	1.8965	1.8965	1.8965	1.8965	1.8965	1.8965
stall time (sec)	4.723	0.041	0.018	0.018	0.018	0.018
elapsed time (sec)	85.835	81.153	81.13	81.13	81.13	81.13
average fetch time (msec)	17.261	18.038	18.066	18.353	18.357	18.126
average disk utilization	0.76	0.42	0.28	0.21	0.17	0.14
Aggressive						
fetches	3943	4797	4976	4943	4863	4856
driver time (sec)	1.9715	2.3985	2.488	2.4715	2.4315	2.428
stall time (sec)	3.995	0.153	0.258	0	0	0.001
elapsed time (sec)	85.182	81.767	81.962	81.687	81.647	81.645
average fetch time (msec)	14.781	16.013	15.626	16.613	16.646	17.033
average disk utilization	0.68	0.47	0.32	0.25	0.2	0.17
Reverse Aggressive						
fetches	3801	3795	3801	3801	3801	3801
driver time (sec)	1.9005	1.8975	1.9005	1.9005	1.9005	1.9005
stall time (sec)	3.775	0.009	0.001	0	0	0.001
elapsed time (sec)	84.891	81.122	81.117	81.116	81.116	81.117
average fetch time (msec)	14.786	17.122	17.184	17.237	17.173	17.464
average disk utilization	0.66	0.4	0.27	0.2	0.16	0.14

Table 33: Performance on the postgres-join trace, cache size 1920.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	3155	3155	3155	3155	3155	3155
driver time (sec)	1.5775	1.5775	1.5775	1.5775	1.5775	1.5775
stall time (sec)	32.755	12.841	6.049	3.226	1.402	0.581
elapsed time (sec)	45.81	25.896	19.104	16.281	14.457	13.636
average fetch time (msec)	14.222	14.786	14.943	15.012	15.279	15.323
average disk utilization	0.98	0.9	0.82	0.73	0.67	0.59
Aggressive						
fetches	3249	3299	3394	3317	3965	4099
driver time (sec)	1.6245	1.6495	1.697	1.6585	1.9825	2.0495
stall time (sec)	32.613	11.717	4.384	1.007	0	0.001
elapsed time (sec)	45.715	24.844	17.559	14.143	13.46	13.528
average fetch time (msec)	13.938	14.118	13.944	14.557	13.923	15.171
average disk utilization	0.99	0.94	0.9	0.85	0.82	0.77
Reverse Aggressive						
fetches	3274	3354	3290	3341	3191	3170
driver time (sec)	1.637	1.677	1.645	1.6705	1.5955	1.585
stall time (sec)	31.026	10.967	4.271	0.353	0.005	0.013
elapsed time (sec)	44.141	24.122	17.394	13.501	13.078	13.076
average fetch time (msec)	13.282	13.614	14.029	13.744	15.048	15.574
average disk utilization	0.99	0.95	0.88	0.85	0.73	0.63

Table 34: Performance on the postgres-select trace, cache size 640.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	32.37	12.647	5.943	3.154	1.402	0.581
elapsed time (sec)	45.39	25.667	18.963	16.174	14.422	13.601
average fetch time (msec)	14.368	14.906	15.044	15.13	15.347	15.413
average disk utilization	0.98	0.9	0.82	0.72	0.66	0.58
Aggressive						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.691	10.772	3.517	0.844	0	0.001
elapsed time (sec)	43.711	23.792	16.537	13.864	13.02	13.021
average fetch time (msec)	13.985	14.173	13.95	14.55	14.446	15.175
average disk utilization	0.99	0.92	0.87	0.81	0.68	0.6
Reverse Aggressive						
fetches	3106	3106	3122	3110	3109	3108
driver time (sec)	1.553	1.553	1.561	1.555	1.5545	1.554
stall time (sec)	28.956	8.461	2.656	0.125	0	0.001
elapsed time (sec)	41.987	21.492	15.695	13.158	13.032	13.033
average fetch time (msec)	13.248	12.668	12.072	13.581	14.419	15.115
average disk utilization	0.98	0.92	0.8	0.8	0.69	0.6

Table 35: Performance on the postgres-select trace, cache size 1920.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	6726	6726	6726	6726	6726	6726
driver time (sec)	3.363	3.363	3.363	3.363	3.363	3.363
stall time (sec)	38.883	6.666	4.75	1.751	0.653	0.128
elapsed time (sec)	72.325	40.108	38.192	35.193	34.095	33.57
average fetch time (msec)	10.533	7.597	14.077	9.972	15.262	10.713
average disk utilization	0.98	0.64	0.83	0.48	0.6	0.36
Aggressive						
fetches	6866	8769	7501	10876	8913	11309
driver time (sec)	3.433	4.3845	3.7505	5.438	4.4565	5.6545
stall time (sec)	38.711	1.996	1.719	0.129	0.134	0.055
elapsed time (sec)	72.223	36.459	35.548	35.646	34.669	35.788
average fetch time (msec)	10.484	7.313	13.856	9.646	15.229	10.67
average disk utilization	1	0.88	0.97	0.74	0.78	0.56
Reverse Aggressive						
fetches	6718	6983	6813	7070	6712	7040
driver time (sec)	3.359	3.4915	3.4065	3.535	3.356	3.52
stall time (sec)	37.569	0.371	1.292	0.096	0.016	0.023
elapsed time (sec)	71.007	33.941	34.777	33.71	33.451	33.622
average fetch time (msec)	10.527	7.58	14.184	9.639	14.91	10.011
average disk utilization	1	0.78	0.93	0.51	0.6	0.35

Table 36: Performance on the xds trace, cache size 640.

Disks	1	2	3	4	5	6
Fixed Horizon						
fetches	5392	5392	5392	5392	5392	5392
driver time (sec)	2.696	2.696	2.696	2.696	2.696	2.696
stall time (sec)	29.047	4.104	2.404	0.892	0.304	0.094
elapsed time (sec)	61.822	36.879	35.179	33.667	33.079	32.869
average fetch time (msec)	10.857	7.848	13.977	10.251	14.806	10.846
average disk utilization	0.95	0.57	0.71	0.41	0.48	0.3
Aggressive						
fetches	5392	7067	5894	8817	7483	9174
driver time (sec)	2.696	3.5335	2.947	4.4085	3.7415	4.587
stall time (sec)	26.626	0.337	0.355	0.129	0.134	0.055
elapsed time (sec)	59.401	33.949	33.381	34.616	33.954	34.721
average fetch time (msec)	10.841	7.613	14.085	10.061	15.434	10.683
average disk utilization	0.98	0.79	0.83	0.64	0.68	0.47
Reverse Aggressive						
fetches	5415	5531	5415	5530	5396	5538
driver time (sec)	2.7075	2.7655	2.7075	2.765	2.698	2.769
stall time (sec)	26.357	0.107	0.206	0.02	0.016	0.017
elapsed time (sec)	59.143	32.951	32.992	32.864	32.793	32.865
average fetch time (msec)	10.828	7.809	14.095	9.927	14.775	10.376
average disk utilization	0.99	0.66	0.77	0.42	0.49	0.29

Table 37: Performance on the xds trace, cache size 1920.

E Performance data: varying *aggressive's* batch size

This section contains the performance data for *aggressive* with varying batch size.

Disks	1	2	3	4	5	6
Batch size 4						
fetches	8812	8813	8812	8812	8812	8813
driver time (sec)	4.406	4.4065	4.406	4.406	4.406	4.4065
stall time (sec)	0.023	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	107.967	107.966	107.963	107.961	107.959	107.958
average fetch time (msec)	3.118	3.134	3.151	3.167	3.18	3.193
average disk utilization	0.25	0.13	0.086	0.065	0.052	0.043
Batch size 8						
fetches	8812	8812	8812	8813	8817	8816
driver time (sec)	4.406	4.406	4.406	4.4065	4.4085	4.408
stall time (sec)	0.021	0.017	0.013	0.009	0.005	0.001
elapsed time (sec)	107.965	107.961	107.957	107.954	107.952	107.947
average fetch time (msec)	3.118	3.135	3.155	3.174	3.182	3.203
average disk utilization	0.25	0.13	0.086	0.065	0.052	0.044
Batch size 16						
fetches	8812	8812	8812	8815	8812	8838
driver time (sec)	4.406	4.406	4.406	4.4075	4.406	4.419
stall time (sec)	0.017	0.009	0.001	0	0	0
elapsed time (sec)	107.961	107.953	107.945	107.946	107.944	107.957
average fetch time (msec)	3.131	3.14	3.162	3.176	3.188	3.206
average disk utilization	0.26	0.13	0.086	0.065	0.052	0.044
Batch size 40						
fetches	8812	8812	8823	8852	8869	8826
driver time (sec)	4.406	4.406	4.4115	4.426	4.4345	4.413
stall time (sec)	0.015	0	0	0	0	0
elapsed time (sec)	107.959	107.944	107.95	107.964	107.973	107.951
average fetch time (msec)	3.141	3.146	3.174	3.198	3.217	3.229
average disk utilization	0.26	0.13	0.086	0.066	0.053	0.044
Batch size 80						
fetches	8812	8813	8812	8853	8883	8812
driver time (sec)	4.406	4.4065	4.406	4.4265	4.4415	4.406
stall time (sec)	0.145	0	0.015	0	0	0
elapsed time (sec)	108.089	107.945	107.959	107.965	107.98	107.944
average fetch time (msec)	3.141	3.148	3.185	3.206	3.221	3.23
average disk utilization	0.26	0.13	0.087	0.066	0.053	0.044
Batch size 160						
fetches	8812	8814	8840	8812	8812	8812
driver time (sec)	4.406	4.407	4.42	4.406	4.406	4.406
stall time (sec)	0.405	0	0.195	0.055	0	0
elapsed time (sec)	108.349	107.945	108.153	107.999	107.944	107.944
average fetch time (msec)	3.15	3.163	3.19	3.206	3.215	3.231
average disk utilization	0.26	0.13	0.087	0.065	0.052	0.044

Table 38: Aggressive performance as a function of batch size on the dinero trace.

Disks	1	2	3	4	5	6
Batch size 4						
fetches	5325	8555	8599	8661	8621	8583
driver time (sec)	2.6625	4.2775	4.2995	4.3305	4.3105	4.2915
stall time (sec)	4.161	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	31.758	29.233	29.253	29.282	29.26	29.239
average fetch time (msec)	5.28	3.349	3.374	3.383	3.366	3.376
average disk utilization	0.89	0.49	0.33	0.25	0.2	0.17
Batch size 8						
fetches	5802	8584	8624	8650	8635	8576
driver time (sec)	2.901	4.292	4.312	4.325	4.3175	4.288
stall time (sec)	1.812	0.017	0.013	0.009	0.005	0.001
elapsed time (sec)	29.647	29.243	29.259	29.268	29.257	29.223
average fetch time (msec)	4.513	3.373	3.379	3.367	3.37	3.356
average disk utilization	0.88	0.5	0.33	0.25	0.2	0.16
Batch size 16						
fetches	6300	8585	8647	8678	8621	8572
driver time (sec)	3.15	4.2925	4.3235	4.339	4.3105	4.286
stall time (sec)	0.881	0.009	0.001	0	0	0
elapsed time (sec)	28.965	29.236	29.259	29.273	29.245	29.22
average fetch time (msec)	3.991	3.416	3.41	3.365	3.39	3.38
average disk utilization	0.87	0.5	0.34	0.25	0.2	0.17
Batch size 40						
fetches	6616	8570	8672	8662	8705	8574
driver time (sec)	3.308	4.285	4.336	4.331	4.3525	4.287
stall time (sec)	0.665	0	0	0	0	0
elapsed time (sec)	28.907	29.219	29.27	29.265	29.287	29.221
average fetch time (msec)	3.796	3.361	3.429	3.381	3.421	3.389
average disk utilization	0.87	0.49	0.34	0.25	0.2	0.17
Batch size 80						
fetches	6931	8570	8677	8747	8713	8572
driver time (sec)	3.4655	4.285	4.3385	4.3735	4.3565	4.286
stall time (sec)	0.911	0	0.007	0	0	0
elapsed time (sec)	29.311	29.219	29.28	29.308	29.291	29.22
average fetch time (msec)	3.758	3.374	3.465	3.423	3.424	3.404
average disk utilization	0.89	0.49	0.34	0.26	0.2	0.17
Batch size 160						
fetches	7360	8699	8662	8678	8621	8572
driver time (sec)	3.68	4.3495	4.331	4.339	4.3105	4.286
stall time (sec)	1.83	0.192	0.187	0	0	0
elapsed time (sec)	30.444	29.476	29.452	29.273	29.245	29.22
average fetch time (msec)	3.757	3.436	3.417	3.413	3.412	3.407
average disk utilization	0.91	0.51	0.33	0.25	0.2	0.17

Table 39: Aggressive performance as a function of batch size on the cscope1 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Batch size 4											
fetches	5982	6098	7029	8240	9574	10678	11406	11717	11619	11102	10662
driver time (sec)	2.991	3.049	3.5145	4.12	4.787	5.339	5.703	5.8585	5.8095	5.551	5.331
stall time (sec)	24.775	9.407	3.736	0.807	0.015	0.013	0.011	0.009	0.005	0.001	0
elapsed time (sec)	64.875	49.565	44.36	42.036	41.911	42.461	42.823	42.977	42.924	42.661	42.44
average fetch time (msec)	10.407	15.336	16.97	17.65	17.787	17.891	17.752	17.551	17.224	17.65	18.201
average disk utilization	0.96	0.94	0.9	0.86	0.81	0.75	0.68	0.6	0.47	0.38	0.29
Batch size 8											
fetches	6009	6194	7277	8541	9952	11014	11587	11758	11565	11134	10662
driver time (sec)	3.0045	3.097	3.6385	4.2705	4.976	5.507	5.7935	5.879	5.7825	5.567	5.331
stall time (sec)	23.367	8.442	3.15	0.299	0.005	0.001	0	0	0	0	0
elapsed time (sec)	63.481	48.648	43.898	41.679	42.09	42.617	42.903	42.988	42.892	42.676	42.44
average fetch time (msec)	10.298	14.97	16.392	17.193	17.527	17.683	17.722	17.532	17.195	17.534	18.135
average disk utilization	0.97	0.95	0.91	0.88	0.83	0.76	0.68	0.6	0.46	0.38	0.28
Batch size 16											
fetches	6044	6321	7583	8956	10299	11364	11617	11758	11565	11134	10662
driver time (sec)	3.022	3.1605	3.7915	4.478	5.1495	5.682	5.8085	5.879	5.7825	5.567	5.331
stall time (sec)	20.041	6.887	2.507	0	0	0	0	0	0	0	0
elapsed time (sec)	60.172	47.157	43.408	41.587	42.259	42.791	42.918	42.988	42.892	42.676	42.44
average fetch time (msec)	9.793	14.201	15.654	16.514	17.138	17.426	17.444	17.343	17.098	17.498	18.085
average disk utilization	0.98	0.95	0.91	0.89	0.84	0.77	0.67	0.59	0.46	0.38	0.28
Batch size 40											
fetches	6171	6592	8208	9684	10892	11553	11728	11884	11654	11164	10662
driver time (sec)	3.0855	3.296	4.104	4.842	5.446	5.7765	5.864	5.942	5.827	5.582	5.331
stall time (sec)	16.349	5.597	1.798	0.156	0.044	0.115	0.06	0.046	0	0	0
elapsed time (sec)	56.544	46.002	43.011	42.107	42.599	43.001	43.033	43.097	42.936	42.691	42.44
average fetch time (msec)	9.099	13.256	14.354	15.55	16.49	17.067	17.266	17.228	16.938	17.369	17.931
average disk utilization	0.99	0.95	0.91	0.89	0.84	0.76	0.67	0.59	0.46	0.38	0.28
Batch size 80											
fetches	6318	7022	8799	10463	11331	11655	11753	11897	11807	11291	10662
driver time (sec)	3.159	3.511	4.3995	5.2315	5.6655	5.8275	5.8765	5.9485	5.9035	5.6455	5.331
stall time (sec)	15.858	4.803	1.017	0.48	0.285	0.4	0.197	0.183	0.091	0.115	0
elapsed time (sec)	56.126	45.423	42.526	42.821	43.06	43.337	43.183	43.241	43.104	42.87	42.44
average fetch time (msec)	8.773	12.278	13.331	14.736	16.237	16.979	17.106	17.009	16.756	17.182	17.717
average disk utilization	0.99	0.95	0.92	0.9	0.85	0.76	0.67	0.58	0.46	0.38	0.28
Batch size 160											
fetches	6771	7778	9478	10967	11325	11942	11859	12039	11619	11130	10604
driver time (sec)	3.3855	3.889	4.739	5.4835	5.6625	5.971	5.9295	6.0195	5.8095	5.565	5.302
stall time (sec)	17.081	5.856	1.078	1.119	1.029	0.956	0.832	0.882	0.559	0.197	0.016
elapsed time (sec)	57.576	46.854	42.926	43.712	43.801	44.036	43.871	44.011	43.478	42.871	42.427
average fetch time (msec)	8.443	11.569	12.688	14.445	16.183	16.833	17.135	17.078	16.679	17.019	17.744
average disk utilization	0.99	0.96	0.93	0.91	0.84	0.76	0.66	0.58	0.45	0.37	0.28

Table 40: Aggressive performance as a function of batch size on the cscope2 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Batch size 4											
fetches	11763	12349	14291	16061	17462	18026	17821	17577	16917	16542	16314
driver time (sec)	5.8815	6.1745	7.1455	8.0305	8.731	9.013	8.9105	8.7885	8.4585	8.271	8.157
stall time (sec)	28.053	6.26	1.518	0.363	0.015	0.013	0.011	0.009	0.005	0.001	0
elapsed time (sec)	108.035	86.535	82.764	82.494	82.847	83.127	83.022	82.898	82.564	82.373	82.258
average fetch time (msec)	8.624	12.843	15.411	16.314	16.698	16.667	16.725	16.905	17.605	17.966	18.49
average disk utilization	0.94	0.92	0.89	0.79	0.7	0.6	0.51	0.45	0.36	0.3	0.23
Batch size 8											
fetches	11779	12627	14881	16441	17635	18081	17894	17565	16902	16546	16284
driver time (sec)	5.8895	6.3135	7.4405	8.2205	8.8175	9.0405	8.947	8.7825	8.451	8.273	8.142
stall time (sec)	25.287	5.205	1.323	0.069	0.005	0.001	0	0	0	0	0
elapsed time (sec)	105.277	85.619	82.864	82.39	82.923	83.142	83.048	82.883	82.552	82.374	82.243
average fetch time (msec)	8.623	12.65	15.121	16.158	16.662	16.568	16.711	16.84	17.523	17.942	18.445
average disk utilization	0.96	0.93	0.91	0.81	0.71	0.6	0.51	0.45	0.36	0.3	0.23
Batch size 16											
fetches	11811	13043	15366	16740	17713	18175	17924	17607	16902	16598	16314
driver time (sec)	5.9055	6.5215	7.683	8.37	8.8565	9.0875	8.962	8.8035	8.451	8.299	8.157
stall time (sec)	21.462	3.911	1.079	0.052	0	0.036	0	0	0	0.033	0
elapsed time (sec)	101.468	84.533	82.863	82.523	82.957	83.224	83.063	82.904	82.552	82.433	82.258
average fetch time (msec)	8.416	12.214	14.719	15.92	16.553	16.541	16.552	16.811	17.462	17.882	18.404
average disk utilization	0.98	0.94	0.91	0.81	0.71	0.6	0.51	0.45	0.36	0.3	0.23
Batch size 40											
fetches	11925	13572	15938	17104	17842	18158	17924	17748	16981	16646	16344
driver time (sec)	5.9625	6.786	7.969	8.552	8.921	9.079	8.962	8.874	8.4905	8.323	8.172
stall time (sec)	15.648	2.862	0.64	0.274	0.264	0.249	0	0.06	0.179	0.204	0
elapsed time (sec)	95.711	83.749	82.71	82.927	83.286	83.429	83.063	83.035	82.77	82.628	82.273
average fetch time (msec)	7.949	11.597	14.215	15.737	16.3	16.35	16.496	16.651	17.391	17.777	18.264
average disk utilization	0.99	0.94	0.91	0.81	0.7	0.59	0.51	0.44	0.36	0.3	0.23
Batch size 80											
fetches	12092	14105	16543	17257	17919	18357	18137	17819	16963	16766	16275
driver time (sec)	6.046	7.0525	8.2715	8.6285	8.9595	9.1785	9.0685	8.9095	8.4815	8.383	8.1375
stall time (sec)	13.943	2.195	0.715	0.584	0.704	0.489	0.219	0.368	0.494	0.338	0.093
elapsed time (sec)	94.09	83.348	83.087	83.313	83.764	83.768	83.388	83.378	83.076	82.822	82.331
average fetch time (msec)	7.741	11.093	13.798	15.546	16.225	16.284	16.339	16.549	17.325	17.639	18.155
average disk utilization	0.99	0.94	0.92	0.81	0.69	0.59	0.51	0.44	0.35	0.3	0.22
Batch size 160											
fetches	12512	14919	16966	17314	18012	18450	18245	17733	16924	16468	16249
driver time (sec)	6.256	7.4595	8.483	8.657	9.006	9.225	9.1225	8.8665	8.462	8.234	8.1245
stall time (sec)	15.216	2.542	1.455	1.297	1.523	0.981	0.913	0.947	0.801	0.373	0.031
elapsed time (sec)	95.573	84.102	84.039	84.055	84.63	84.307	84.136	83.914	83.364	82.708	82.256
average fetch time (msec)	7.512	10.643	13.615	15.645	16.247	16.196	16.337	16.62	17.185	17.6	18.164
average disk utilization	0.98	0.94	0.92	0.81	0.69	0.59	0.51	0.44	0.35	0.29	0.22

Table 41: Aggressive performance as a function of batch size on the cscope3 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Batch size 4											
fetches	6520	6597	6803	7178	7957	8946	10017	10992	12009	11530	11315
driver time (sec)	3.26	3.2985	3.4015	3.589	3.9785	4.473	5.0085	5.496	6.0045	5.765	5.6575
stall time (sec)	62.407	24.43	10.612	3.948	2.034	0.371	0.011	0.009	0.005	0.001	0
elapsed time (sec)	104.383	66.445	52.73	46.253	44.729	43.56	43.736	44.221	44.726	44.482	44.374
average fetch time (msec)	13.918	15.694	16.698	17.342	17.959	17.334	17.074	16.896	16.137	15.917	16.198
average disk utilization	0.87	0.78	0.72	0.67	0.64	0.59	0.56	0.52	0.43	0.34	0.26
Batch size 8											
fetches	6532	6636	6883	7349	8475	9376	10423	11296	12085	11709	11457
driver time (sec)	3.266	3.318	3.4415	3.6745	4.2375	4.688	5.2115	5.648	6.0425	5.8545	5.7285
stall time (sec)	60.737	22.944	9.255	3.264	1.495	0.035	0	0	0	0	0
elapsed time (sec)	102.719	64.978	51.413	45.655	44.449	43.439	43.928	44.364	44.759	44.571	44.445
average fetch time (msec)	13.732	15.366	16.344	16.857	16.92	16.836	16.79	16.886	15.76	15.862	16.169
average disk utilization	0.87	0.78	0.73	0.68	0.65	0.61	0.57	0.54	0.43	0.35	0.26
Batch size 16											
fetches	6553	6631	6959	7551	8908	9975	10860	11795	12487	11980	11499
driver time (sec)	3.2765	3.3155	3.4795	3.7755	4.454	4.9875	5.43	5.8975	6.2435	5.99	5.7495
stall time (sec)	59.057	20.883	7.529	2.495	0.826	0	0	0	0	0	0
elapsed time (sec)	101.05	62.915	49.725	44.987	43.996	43.704	44.146	44.614	44.96	44.706	44.466
average fetch time (msec)	13.523	15.002	15.658	16.247	15.973	16.255	16.513	16.529	15.635	15.695	16.172
average disk utilization	0.88	0.79	0.73	0.68	0.65	0.62	0.58	0.55	0.43	0.35	0.26
Batch size 40											
fetches	6601	6888	7287	8524	9998	10670	11662	12237	12721	12148	11517
driver time (sec)	3.3005	3.444	3.6435	4.262	4.999	5.335	5.831	6.1185	6.3605	6.074	5.7585
stall time (sec)	56.841	18.58	6.384	1.608	0	0	0	0	0	0	0
elapsed time (sec)	98.858	60.74	48.744	44.586	43.715	44.051	44.547	44.835	45.077	44.79	44.475
average fetch time (msec)	13.225	14.259	14.645	14.741	14.873	15.694	16.011	15.98	15.585	15.677	15.951
average disk utilization	0.88	0.81	0.73	0.7	0.68	0.63	0.6	0.55	0.44	0.35	0.26
Batch size 80											
fetches	6690	7128	7930	9430	11022	11778	12142	12540	12643	12026	11487
driver time (sec)	3.345	3.564	3.965	4.715	5.511	5.889	6.071	6.27	6.3215	6.013	5.7435
stall time (sec)	54.58	16.457	5.461	0.836	0	0.021	0	0.001	0	0	0
elapsed time (sec)	96.641	58.737	48.142	44.267	44.227	44.626	44.787	44.987	45.038	44.729	44.46
average fetch time (msec)	12.889	13.571	13.547	13.629	14.129	15.243	15.763	15.963	15.638	15.547	15.83
average disk utilization	0.89	0.82	0.74	0.73	0.7	0.67	0.61	0.56	0.44	0.35	0.26
Batch size 160											
fetches	7007	7690	9473	10657	11387	11825	12417	12521	12885	12062	11487
driver time (sec)	3.5035	3.845	4.7365	5.3285	5.6935	5.9125	6.2085	6.2605	6.4425	6.031	5.7435
stall time (sec)	51.598	17.187	5.473	1.041	0.7	0.575	0.27	0.37	0	0	0
elapsed time (sec)	93.818	59.748	48.926	45.086	45.11	45.204	45.195	45.347	45.159	44.747	44.46
average fetch time (msec)	12.508	12.9	12.754	13.102	14.031	15.132	15.654	15.918	15.514	15.515	15.792
average disk utilization	0.93	0.83	0.82	0.77	0.71	0.66	0.61	0.55	0.44	0.35	0.26

Table 42: Aggressive performance as a function of batch size on the glimpse trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Batch size 4											
fetches	2885	2909	2916	2952	3048	3332	3586	3779	4091	4285	4651
driver time (sec)	1.4425	1.4545	1.458	1.476	1.524	1.666	1.793	1.8895	2.0455	2.1425	2.3255
stall time (sec)	16.476	8.221	5.117	2.955	1.414	0.737	0.174	0.009	0.005	0.001	0
elapsed time (sec)	26.083	17.84	14.74	12.596	11.103	10.568	10.132	10.063	10.215	10.308	10.49
average fetch time (msec)	8.679	11.449	13.866	15.474	16.32	16.472	17.13	17.175	17.62	18.017	18.261
average disk utilization	0.96	0.93	0.91	0.91	0.9	0.87	0.87	0.81	0.71	0.62	0.51
Batch size 8											
fetches	2892	2918	2942	3021	3192	3505	3734	3951	4183	4410	4687
driver time (sec)	1.446	1.459	1.471	1.5105	1.596	1.7525	1.867	1.9755	2.0915	2.205	2.3435
stall time (sec)	17.041	8.126	4.675	2.65	1.241	0.265	0.023	0.012	0	0	0
elapsed time (sec)	26.652	17.75	14.311	12.325	11.002	10.182	10.055	10.152	10.256	10.37	10.508
average fetch time (msec)	8.981	11.609	13.686	14.984	15.797	15.958	16.444	16.804	17.436	17.624	18.105
average disk utilization	0.97	0.95	0.94	0.92	0.92	0.92	0.87	0.82	0.71	0.62	0.5
Batch size 16											
fetches	2896	2942	2982	3102	3310	3626	3856	4107	4329	4559	4748
driver time (sec)	1.448	1.471	1.491	1.551	1.655	1.813	1.928	2.0535	2.1645	2.2795	2.374
stall time (sec)	16.552	7.34	3.845	2.052	0.579	0.287	0	0	0	0	0
elapsed time (sec)	26.165	16.976	13.501	11.768	10.399	10.265	10.093	10.218	10.329	10.444	10.539
average fetch time (msec)	8.919	11.194	12.99	14.199	14.932	15.516	15.914	16.277	17.002	17.214	17.786
average disk utilization	0.99	0.97	0.96	0.94	0.95	0.91	0.87	0.82	0.71	0.63	0.5
Batch size 40											
fetches	2934	2982	3137	3297	3560	3893	4105	4338	4610	4776	4741
driver time (sec)	1.467	1.491	1.5685	1.6485	1.78	1.9465	2.0525	2.169	2.305	2.388	2.3705
stall time (sec)	15.58	6.329	3.433	1.594	0.368	0.22	0.077	0.013	0	0	0
elapsed time (sec)	25.212	15.985	13.166	11.407	10.313	10.331	10.294	10.347	10.47	10.553	10.535
average fetch time (msec)	8.523	10.583	12.037	13.131	13.917	14.641	15.208	15.687	16.162	16.703	17.38
average disk utilization	0.99	0.99	0.96	0.95	0.96	0.92	0.87	0.82	0.71	0.63	0.49
Batch size 80											
fetches	2981	3142	3257	3571	3909	4326	4496	4692	4941	4862	4695
driver time (sec)	1.4905	1.571	1.6285	1.7855	1.9545	2.163	2.248	2.346	2.4705	2.431	2.3475
stall time (sec)	15.245	6.26	3.176	1.554	0.798	0.392	0.244	0.332	0.05	0.083	0
elapsed time (sec)	24.9	15.996	12.969	11.504	10.917	10.72	10.657	10.843	10.685	10.679	10.512
average fetch time (msec)	8.248	9.932	11.459	12.33	13.341	13.85	14.594	15.099	15.92	16.378	16.986
average disk utilization	0.99	0.98	0.96	0.96	0.96	0.93	0.88	0.82	0.74	0.62	0.47
Batch size 160											
fetches	3134	3356	3823	4207	4560	4644	4914	4721	4788	4688	4630
driver time (sec)	1.567	1.678	1.9115	2.1035	2.28	2.322	2.457	2.3605	2.394	2.344	2.315
stall time (sec)	15.141	6.183	3.941	2.35	1.72	1.466	0.879	0.953	0.441	0.386	0
elapsed time (sec)	24.873	16.026	14.017	12.618	12.165	11.953	11.501	11.478	11	10.895	10.48
average fetch time (msec)	7.856	9.362	10.522	11.344	12.389	13.653	14.519	15.356	15.746	16.049	17.023
average disk utilization	0.99	0.98	0.96	0.95	0.93	0.88	0.89	0.79	0.69	0.58	0.47

Table 43: Aggressive performance as a function of batch size on the ld trace.

Disks	1	2	3	4	5	6
Batch size 4						
fetches	3925	5451	6044	6062	5978	5925
driver time (sec)	1.9625	2.7255	3.022	3.031	2.989	2.9625
stall time (sec)	11.351	0.021	0.019	0.017	0.015	0.013
elapsed time (sec)	92.529	81.962	82.257	82.264	82.22	82.191
average fetch time (msec)	18.345	17.669	16.975	17.051	17.256	17.288
average disk utilization	0.78	0.59	0.42	0.31	0.25	0.21
Batch size 8						
fetches	4051	5642	6116	6078	5998	5919
driver time (sec)	2.0255	2.821	3.058	3.039	2.999	2.9595
stall time (sec)	9.318	0.017	0.013	0.009	0.005	0.001
elapsed time (sec)	90.559	82.054	82.287	82.264	82.22	82.176
average fetch time (msec)	17.772	17.483	16.701	16.848	17.034	17.102
average disk utilization	0.79	0.6	0.41	0.31	0.25	0.21
Batch size 16						
fetches	4233	5718	6170	6156	6047	5951
driver time (sec)	2.1165	2.859	3.085	3.078	3.0235	2.9755
stall time (sec)	7.437	0.009	0.001	0	0	0
elapsed time (sec)	88.769	82.084	82.302	82.294	82.239	82.191
average fetch time (msec)	16.928	17.15	16.438	16.578	16.706	16.837
average disk utilization	0.81	0.6	0.41	0.31	0.25	0.2
Batch size 40						
fetches	4510	5836	6225	6239	6137	5954
driver time (sec)	2.255	2.918	3.1125	3.1195	3.0685	2.977
stall time (sec)	4.611	0.152	0.258	0.258	0	0.207
elapsed time (sec)	86.082	82.286	82.586	82.593	82.284	82.4
average fetch time (msec)	15.643	16.576	15.929	16.083	16.36	16.349
average disk utilization	0.82	0.59	0.4	0.3	0.24	0.2
Batch size 80						
fetches	4698	5900	6307	6279	6140	5992
driver time (sec)	2.349	2.95	3.1535	3.1395	3.07	2.996
stall time (sec)	3.994	0.712	0.758	0.673	0.129	0.487
elapsed time (sec)	85.559	82.878	83.127	83.028	82.415	82.699
average fetch time (msec)	15.032	16.431	15.682	15.941	16.042	16.047
average disk utilization	0.83	0.58	0.4	0.3	0.24	0.19
Batch size 160						
fetches	4890	5940	6387	6270	6090	6176
driver time (sec)	2.445	2.97	3.1935	3.135	3.045	3.088
stall time (sec)	3.774	1.591	1.522	1.232	0.424	0.922
elapsed time (sec)	85.435	83.777	83.931	83.583	82.685	83.226
average fetch time (msec)	14.506	16.232	15.549	15.557	15.763	15.713
average disk utilization	0.83	0.58	0.39	0.29	0.23	0.19

Table 44: Aggressive performance as a function of batch size on the postgres-join trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Batch size 4											
fetches	3085	3085	3085	3085	3085	3220	3633	3937	3902	3852	3731
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.61	1.8165	1.9685	1.951	1.926	1.8655
stall time (sec)	39.507	13.89	5.578	1.786	0.177	0.013	0.011	0.009	0.005	0.001	0
elapsed time (sec)	52.527	26.91	18.598	14.806	13.197	13.101	13.305	13.455	13.434	13.405	13.343
average fetch time (msec)	16.582	16.121	15.884	15.833	15.544	15.172	15.412	15.274	14.797	14.8	15.155
average disk utilization	0.97	0.92	0.88	0.82	0.73	0.62	0.6	0.56	0.43	0.35	0.26
Batch size 8											
fetches	3085	3085	3085	3085	3085	3317	3826	4068	3975	3854	3731
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.6585	1.913	2.034	1.9875	1.927	1.8655
stall time (sec)	37.572	13.285	4.968	1.198	0.005	0.001	0	0	0	0	0
elapsed time (sec)	50.592	26.305	17.988	14.218	13.025	13.137	13.391	13.512	13.465	13.405	13.343
average fetch time (msec)	15.972	15.679	15.45	15.156	15.113	15.036	15.221	14.961	14.501	14.653	14.895
average disk utilization	0.97	0.92	0.88	0.82	0.72	0.63	0.62	0.56	0.43	0.35	0.26
Batch size 16											
fetches	3085	3085	3085	3085	3286	3563	3989	4158	3975	3854	3731
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.643	1.7815	1.9945	2.079	1.9875	1.927	1.8655
stall time (sec)	35.953	12.022	4.149	0.844	0	0	0	0	0	0	0
elapsed time (sec)	48.973	25.042	17.169	13.864	13.121	13.259	13.472	13.557	13.465	13.405	13.343
average fetch time (msec)	15.468	15.047	14.865	14.55	13.923	14.605	14.883	14.47	14.155	14.44	14.675
average disk utilization	0.97	0.93	0.89	0.81	0.7	0.65	0.63	0.55	0.42	0.35	0.26
Batch size 40											
fetches	3085	3085	3085	3085	3498	3947	4207	4218	4059	3933	3741
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.749	1.9735	2.1035	2.109	2.0295	1.9665	1.8705
stall time (sec)	33.015	10.772	3.517	0.11	0.008	0.133	0.185	0.13	0	0.038	0
elapsed time (sec)	46.035	23.792	16.537	13.13	13.235	13.584	13.766	13.717	13.507	13.482	13.348
average fetch time (msec)	14.695	14.173	13.95	13.619	13.303	14.127	14.887	14.086	13.765	13.766	13.85
average disk utilization	0.98	0.92	0.87	0.8	0.7	0.68	0.65	0.54	0.41	0.33	0.24
Batch size 80											
fetches	3085	3085	3085	3222	3873	4118	4331	4157	4061	4100	3741
driver time (sec)	1.5425	1.5425	1.5425	1.611	1.9365	2.059	2.1655	2.0785	2.0305	2.05	1.8705
stall time (sec)	30.691	9.611	2.918	0.584	0.358	0.568	0.57	0.439	0	0.147	0
elapsed time (sec)	43.711	22.631	15.938	13.673	13.772	14.105	14.213	13.995	13.508	13.675	13.348
average fetch time (msec)	13.985	13.567	13.137	12.899	13.015	14.239	14.299	13.683	13.409	13.434	13.109
average disk utilization	0.99	0.92	0.85	0.76	0.73	0.69	0.62	0.51	0.4	0.34	0.23
Batch size 160											
fetches	3085	3085	3388	3762	4126	4355	4340	4237	3933	3820	3681
driver time (sec)	1.5425	1.5425	1.694	1.881	2.063	2.1775	2.17	2.1185	1.9665	1.91	1.8405
stall time (sec)	28.957	8.422	3.301	1.369	0.888	1.044	0.968	0.831	0.048	0.182	0
elapsed time (sec)	41.977	21.442	16.473	14.728	14.429	14.699	14.616	14.427	13.492	13.57	13.318
average fetch time (msec)	13.296	12.704	12.173	12.398	12.972	13.751	13.798	12.985	12.565	12.885	12.883
average disk utilization	0.98	0.91	0.83	0.79	0.74	0.68	0.59	0.48	0.37	0.3	0.22

Table 45: Aggressive performance as a function of batch size on the postgres-select trace.

Disks	1	2	3	4	5	6
Batch size 4						
fetches	5858	6939	6023	9564	8040	10098
driver time (sec)	2.929	3.4695	3.0115	4.782	4.02	5.049
stall time (sec)	32.888	0.162	1.06	0.02	0.018	0.015
elapsed time (sec)	65.896	33.71	34.15	34.881	34.117	35.143
average fetch time (msec)	10.977	7.635	14.623	9.75	15.515	10.778
average disk utilization	0.98	0.79	0.86	0.67	0.73	0.52
Batch size 8						
fetches	5862	7200	6202	9827	8068	10215
driver time (sec)	2.931	3.6	3.101	4.9135	4.034	5.1075
stall time (sec)	32.197	0.077	0.59	0.071	0.065	0.055
elapsed time (sec)	65.207	33.756	33.77	35.063	34.178	35.241
average fetch time (msec)	10.914	7.583	14.355	9.788	15.45	10.711
average disk utilization	0.98	0.81	0.88	0.69	0.73	0.52
Batch size 16						
fetches	5868	7522	6306	9831	8312	10124
driver time (sec)	2.934	3.761	3.153	4.9155	4.156	5.062
stall time (sec)	31.504	0.192	0.205	0.129	0.133	0.076
elapsed time (sec)	64.517	34.032	33.437	35.123	34.368	35.217
average fetch time (msec)	10.854	7.564	14.151	9.801	15.454	10.6
average disk utilization	0.99	0.84	0.89	0.69	0.75	0.51
Batch size 40						
fetches	5890	7778	6563	9929	8418	10353
driver time (sec)	2.945	3.889	3.2815	4.9645	4.209	5.1765
stall time (sec)	30.61	0.337	0.356	0.232	0.312	0.18
elapsed time (sec)	63.634	34.305	33.716	35.275	34.6	35.435
average fetch time (msec)	10.745	7.496	14.101	9.92	15.441	10.63
average disk utilization	0.99	0.85	0.91	0.7	0.75	0.52
Batch size 80						
fetches	5925	8126	6838	10150	8789	10461
driver time (sec)	2.9625	4.063	3.419	5.075	4.3945	5.2305
stall time (sec)	30.667	0.507	0.613	0.399	0.582	0.525
elapsed time (sec)	63.708	34.649	34.111	35.553	35.055	35.834
average fetch time (msec)	10.711	7.386	14.051	9.957	15.258	10.584
average disk utilization	1	0.87	0.94	0.71	0.77	0.51
Batch size 160						
fetches	6005	8198	7519	10300	8601	10488
driver time (sec)	3.0025	4.099	3.7595	5.15	4.3005	5.244
stall time (sec)	31.271	0.951	1.998	1.331	1.608	0.978
elapsed time (sec)	64.352	35.129	35.836	36.56	35.987	36.301
average fetch time (msec)	10.707	7.447	13.692	9.715	15.213	10.363
average disk utilization	1	0.87	0.96	0.68	0.73	0.5

Table 46: Aggressive performance as a function of batch size on the xds trace.

F Performance data: varying *reverse aggressive*'s parameters

This section contains the performance data for *reverse aggressive* with varying batch sizes and fetch time estimates. For brevity, only the elapsed times are shown.

Disks	1	2	3	4	5	6
Fetch time 4						
Batch size 4	105.932	106.013	106.992	107.961	107.959	107.957
Batch size 8	105.946	106.013	106.991	107.953	107.949	107.945
Batch size 16	105.976	106.093	106.988	107.944	107.944	107.944
Batch size 40	106.03	106.348	107.021	107.944	107.944	107.944
Batch size 80	106.283	106.798	107.468	107.944	107.944	107.947
Batch size 160	106.825	107.719	108.033	108.044	107.944	107.954
Fetch time 8						
Batch size 4	105.931	105.949	105.979	106.013	106.502	106.99
Batch size 8	105.944	105.981	106.023	106.119	106.502	106.99
Batch size 16	105.972	106.094	106.13	106.335	106.512	107.004
Batch size 40	106.01	106.332	106.476	106.891	107.282	107.542
Batch size 80	106.22	106.716	107.104	107.805	107.944	107.944
Batch size 160	106.708	107.425	108.11	108.148	107.944	107.954
Fetch time 16						
Batch size 4	105.929	105.946	105.978	105.986	106.047	106.06
Batch size 8	105.945	105.977	106.028	106.09	106.146	106.191
Batch size 16	105.976	106.093	106.156	106.233	106.344	106.459
Batch size 40	105.975	106.32	106.478	106.743	107.241	107.499
Batch size 80	106.181	106.716	107.227	107.685	107.939	107.944
Batch size 160	106.684	107.425	108.149	108.059	107.944	108.014
Fetch time 32						
Batch size 4	105.927	105.945	105.978	105.981	106.047	106.069
Batch size 8	105.942	105.977	106.064	106.091	106.134	106.163
Batch size 16	105.974	106.093	106.161	106.253	106.329	106.402
Batch size 40	105.982	106.288	106.508	106.783	107.107	107.454
Batch size 80	106.15	106.716	107.371	107.659	107.935	107.948
Batch size 160	106.612	107.398	108.159	108.074	107.944	108.014
Fetch time 64						
Batch size 4	105.927	105.941	105.972	105.978	106.047	106.063
Batch size 8	105.941	105.977	106.025	106.106	106.139	106.171
Batch size 16	105.969	106.089	106.17	106.203	106.302	106.369
Batch size 40	105.987	106.304	106.464	106.749	107.011	107.513
Batch size 80	106.15	106.716	107.268	107.594	107.907	107.944
Batch size 160	106.628	107.407	108.153	108.163	107.944	107.944
Fetch time 128						
Batch size 4	105.927	105.941	105.972	105.97	106.01	106.063
Batch size 8	105.941	105.969	106.017	106.09	106.135	106.171
Batch size 16	105.969	106.089	106.17	106.336	106.314	106.419
Batch size 40	105.987	106.312	106.539	106.689	106.994	107.484
Batch size 80	106.15	106.716	107.27	107.603	107.932	107.952
Batch size 160	106.66	107.425	108.154	108.164	107.944	107.984

Table 47: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the dinero trace.

Disks	1	2	3	4	5	6
Fetch time 4						
Batch size 4	29.884	27.485	28.328	29.219	29.217	29.215
Batch size 8	29.509	27.529	28.328	29.211	29.207	29.203
Batch size 16	29.42	27.589	28.327	29.202	29.202	29.202
Batch size 40	29.339	27.807	28.351	29.202	29.202	29.202
Batch size 80	29.098	28.238	28.778	29.202	29.202	29.202
Batch size 160	29.894	29.255	29.404	29.348	29.202	29.202
Fetch time 8						
Batch size 4	30.199	27.47	27.489	27.526	27.854	28.326
Batch size 8	30.072	27.486	27.532	27.593	27.851	28.326
Batch size 16	29.987	27.536	27.619	27.757	27.872	28.349
Batch size 40	29.479	27.678	27.927	28.231	28.528	28.937
Batch size 80	28.921	28.015	28.522	29.065	29.195	29.202
Batch size 160	29.792	29.038	29.351	29.438	29.202	29.202
Fetch time 16						
Batch size 4	30.379	27.461	27.477	27.498	27.515	27.541
Batch size 8	30.34	27.51	27.517	27.556	27.603	27.661
Batch size 16	30.177	27.525	27.639	27.682	27.78	27.894
Batch size 40	29.683	27.756	27.869	28.066	28.362	28.664
Batch size 80	29.105	28.104	28.463	28.906	29.181	29.202
Batch size 160	30.051	29.045	29.377	29.318	29.202	29.202
Fetch time 32						
Batch size 4	30.499	27.457	27.471	27.513	27.528	27.527
Batch size 8	30.423	27.507	27.513	27.544	27.592	27.673
Batch size 16	30.351	27.51	27.605	27.673	27.767	27.888
Batch size 40	30.048	27.725	27.954	28.159	28.374	28.675
Batch size 80	29.672	28.072	28.618	28.907	29.188	29.202
Batch size 160	30.773	29.036	29.393	29.455	29.202	29.202
Fetch time 64						
Batch size 4	30.544	27.453	27.471	27.506	27.518	27.515
Batch size 8	30.465	27.48	27.513	27.563	27.598	27.662
Batch size 16	30.319	27.513	27.61	27.654	27.743	27.855
Batch size 40	30.171	27.74	27.91	28.054	28.459	28.722
Batch size 80	30.055	28.087	28.521	28.904	29.185	29.202
Batch size 160	30.259	29.037	29.295	29.273	29.202	29.202
Fetch time 128						
Batch size 4	30.559	27.453	27.465	27.505	27.517	27.515
Batch size 8	30.521	27.48	27.506	27.548	27.598	27.657
Batch size 16	30.433	27.513	27.608	27.745	27.753	27.874
Batch size 40	30.224	27.752	27.998	28.075	28.4	28.759
Batch size 80	30.088	28.111	28.529	28.898	29.189	29.202
Batch size 160	30.169	29.053	29.324	29.236	29.202	29.202

Table 48: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the cscope1 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 4											
Batch size 4	73.646	76.154	58.632	46.844	42.027	42.025	42.023	42.021	42.017	42.028	42.02
Batch size 8	68.713	73.32	56.656	45.18	42.017	42.013	42.012	42.02	42.012	42.019	42.018
Batch size 16	64.597	69.609	54.03	42.527	42.012	42.012	42.012	42.012	42.012	42.012	42.014
Batch size 40	60.204	65.388	50.161	42.228	42.056	42.167	42.132	42.118	42.012	42.167	42.063
Batch size 80	58.676	61.528	46.006	42.612	42.377	42.467	42.295	42.303	42.242	42.397	42.414
Batch size 160	58.824	58.068	43.961	43.221	43.178	43.163	43.107	43.137	43.002	42.705	42.431
Fetch time 8											
Batch size 4	66.301	53.815	49.939	45.845	42.027	42.025	42.023	42.021	42.017	42.028	42.02
Batch size 8	65.976	52.196	47.404	44.208	42.017	42.013	42.012	42.02	42.012	42.019	42.018
Batch size 16	63.699	50.01	44.751	42.175	42.012	42.012	42.012	42.012	42.012	42.012	42.014
Batch size 40	61.436	49.007	43.106	42.227	42.056	42.167	42.132	42.118	42.012	42.167	42.063
Batch size 80	59.443	48.204	42.297	42.612	42.377	42.467	42.295	42.303	42.242	42.397	42.414
Batch size 160	59.338	49.797	42.507	43.221	43.178	43.163	43.108	43.137	43.002	42.705	42.431
Fetch time 16											
Batch size 4	65.726	55.805	46.45	41.044	40.724	41.196	41.659	41.944	42.014	42.028	42.02
Batch size 8	66.026	53.985	45.297	40.555	40.723	41.195	41.666	41.983	42.012	42.019	42.018
Batch size 16	64.078	52.309	44.325	40.254	40.734	41.215	41.694	42.005	42.012	42.012	42.014
Batch size 40	61.326	51.027	43.843	40.661	40.83	41.439	41.892	42.118	42.012	42.167	42.063
Batch size 80	59.733	48.032	42.27	41.451	41.396	41.849	42.191	42.303	42.242	42.397	42.413
Batch size 160	58.255	48.498	42.151	42.79	43.061	43.116	43.106	43.123	43.002	42.705	42.431
Fetch time 32											
Batch size 4	66.369	55.811	47.22	42.32	40.412	40.16	40.163	40.251	40.722	41.209	41.918
Batch size 8	66.538	54.272	46.13	41.626	40.176	40.197	40.218	40.265	40.734	41.218	41.953
Batch size 16	64.857	52.296	44.816	40.717	40.258	40.309	40.344	40.375	40.77	41.255	41.994
Batch size 40	61.57	50.761	43.215	40.604	40.542	40.8	40.815	40.925	41.024	41.565	42.063
Batch size 80	60.573	48.144	41.63	41.277	41.251	41.61	41.587	41.846	42.048	42.355	42.415
Batch size 160	59.52	47.419	42.172	42.502	42.862	43.117	42.985	42.976	43.202	43.14	42.431
Fetch time 64											
Batch size 4	66.686	55.88	47.224	42.28	40.322	40.162	40.164	40.176	40.18	40.214	40.265
Batch size 8	66.943	54.08	46.155	41.415	40.185	40.191	40.201	40.251	40.263	40.312	40.39
Batch size 16	65.161	52.358	44.88	40.69	40.255	40.279	40.321	40.363	40.436	40.503	40.655
Batch size 40	62.331	50.149	43.592	40.636	40.549	40.706	40.763	40.865	40.942	41.286	41.539
Batch size 80	60.58	47.05	41.959	41.281	41.211	41.485	41.504	41.703	41.978	42.323	42.296
Batch size 160	59.045	47.689	42.042	42.517	42.849	43.065	43.107	43.159	43.189	43.294	42.313
Fetch time 128											
Batch size 4	66.981	56.125	47.224	42.265	40.322	40.158	40.164	40.177	40.19	40.229	40.236
Batch size 8	67.219	54.329	46.125	41.412	40.187	40.196	40.218	40.237	40.255	40.298	40.367
Batch size 16	65.627	52.364	44.691	40.759	40.244	40.267	40.361	40.346	40.394	40.479	40.617
Batch size 40	61.941	50.631	43.293	40.592	40.501	40.68	40.833	40.915	40.876	41.199	41.475
Batch size 80	60.618	48.25	41.506	41.355	41.268	41.476	41.473	41.665	41.959	42.289	42.339
Batch size 160	59.294	46.826	41.969	42.696	42.856	43.096	43.044	42.959	43.168	43.314	42.311

Table 49: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the cscope2 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 4											
Batch size 4	120.374	106.657	86.841	82.28	81.905	81.903	81.901	81.899	81.895	81.906	81.898
Batch size 8	116.187	102.762	83.583	82.011	81.895	81.891	81.89	81.898	81.89	81.897	81.896
Batch size 16	112.783	96.87	82.994	81.957	81.89	81.926	81.89	81.89	81.89	81.968	81.892
Batch size 40	107.785	91.908	82.549	82.224	82.184	82.154	81.964	81.993	82.133	82.226	81.999
Batch size 80	105.834	87.809	82.778	82.593	82.617	82.408	82.259	82.375	82.433	82.482	82.347
Batch size 160	104.065	85.424	83.448	83.379	83.458	83.014	83.052	83.083	82.751	82.653	82.354
Fetch time 8											
Batch size 4	118.714	99.498	82.058	82.09	81.905	81.903	81.901	81.899	81.895	81.906	81.898
Batch size 8	117.231	96.441	81.642	81.832	81.895	81.891	81.89	81.898	81.89	81.897	81.896
Batch size 16	114.398	93.24	81.419	81.797	81.89	81.926	81.89	81.89	81.89	81.968	81.892
Batch size 40	109.654	90.407	81.011	82.111	82.184	82.154	81.964	81.993	82.133	82.226	81.999
Batch size 80	105.912	87.902	81.411	82.556	82.617	82.408	82.259	82.375	82.433	82.482	82.347
Batch size 160	105.316	85.933	82.611	83.379	83.457	83.014	83.052	83.083	82.751	82.653	82.354
Fetch time 16											
Batch size 4	120.959	99.258	86.509	82.047	80.095	80.317	81.024	81.7	81.895	81.906	81.898
Batch size 8	119.269	96.896	85.597	81.219	80.072	80.316	81.025	81.72	81.89	81.897	81.896
Batch size 16	116.419	92.958	84.417	80.524	80.17	80.371	81.047	81.738	81.89	81.968	81.892
Batch size 40	111.804	89.826	82.373	80.647	80.737	80.806	81.203	81.91	82.133	82.226	81.999
Batch size 80	107.929	86.972	81.381	81.365	81.597	81.633	81.846	82.362	82.431	82.482	82.347
Batch size 160	106.978	84.24	82.473	82.839	83.309	82.926	83.044	83.094	82.751	82.653	82.354
Fetch time 32											
Batch size 4	121.953	100.686	86.669	82.028	80.113	80.043	80.05	80.058	80.082	80.33	81.644
Batch size 8	121.183	98.076	85.666	81.279	80.056	80.08	80.091	80.127	80.168	80.341	81.668
Batch size 16	118.298	94.376	84.202	80.741	80.13	80.202	80.214	80.259	80.357	80.545	81.709
Batch size 40	112.982	90.68	82.261	80.612	80.66	80.742	80.645	80.757	81.137	81.409	81.92
Batch size 80	109.193	86.671	81.336	81.293	81.505	81.491	81.592	81.856	82.252	82.319	82.295
Batch size 160	108.381	84.039	82.41	82.678	83.19	82.925	83.065	83.095	82.857	82.675	82.324
Fetch time 64											
Batch size 4	122.666	101.556	87.388	81.982	80.065	80.035	80.038	80.051	80.065	80.094	80.131
Batch size 8	121.796	98.68	86.101	81.224	80.047	80.07	80.078	80.117	80.145	80.185	80.266
Batch size 16	118.707	94.765	84.931	80.619	80.111	80.186	80.195	80.248	80.319	80.452	80.535
Batch size 40	113.545	91.279	82.495	80.602	80.615	80.661	80.584	80.723	81.03	81.328	81.463
Batch size 80	108.795	86.447	81.422	81.186	81.401	81.361	81.526	81.765	82.225	82.32	82.2
Batch size 160	106.666	84.729	82.374	82.613	83.161	82.904	82.979	83.022	82.669	82.372	82.294
Fetch time 128											
Batch size 4	123.14	102.073	87.561	81.879	80.069	80.032	80.054	80.051	80.065	80.109	80.111
Batch size 8	122.151	98.879	86	81.045	80.049	80.07	80.084	80.111	80.12	80.174	80.244
Batch size 16	118.865	94.58	84.897	80.559	80.111	80.176	80.177	80.207	80.274	80.445	80.495
Batch size 40	113.657	91.243	82.453	80.625	80.587	80.641	80.562	80.67	80.96	81.225	81.426
Batch size 80	109.424	87.42	81.338	81.213	81.425	81.377	81.401	81.753	82.214	82.361	82.153
Batch size 160	107.025	84.923	82.402	82.926	83.243	82.917	83.057	83.083	82.857	82.548	82.339

Table 50: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the cscope3 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 4											
Batch size 4	120.334	85.098	65.108	53.355	47.429	44.047	43.626	43.624	43.62	43.62	43.615
Batch size 8	117.591	83.434	63.074	51.409	46.198	43.674	43.615	43.623	43.615	43.617	43.615
Batch size 16	114.802	81.065	60.735	49.93	44.488	43.615	43.615	43.615	43.615	43.615	43.615
Batch size 40	111.624	76.967	57.331	46.21	44.075	43.615	43.615	43.615	43.615	43.615	43.615
Batch size 80	109.285	71.402	53.632	44.116	43.615	43.665	43.649	43.735	43.615	43.794	43.665
Batch size 160	104.802	67.211	51.041	44.23	44.131	44.265	44.2	44.229	43.615	43.979	43.655
Fetch time 8											
Batch size 4	107.913	73.03	64.973	53.265	47.429	44.047	43.626	43.624	43.62	43.62	43.615
Batch size 8	105.957	71.325	62.994	51.409	46.198	43.674	43.615	43.623	43.615	43.617	43.615
Batch size 16	103.948	68.957	60.837	49.93	44.488	43.615	43.615	43.615	43.615	43.615	43.615
Batch size 40	101.399	65.479	57.211	46.21	44.075	43.615	43.615	43.615	43.615	43.615	43.615
Batch size 80	99.3	60.248	53.496	44.116	43.615	43.665	43.649	43.735	43.615	43.794	43.665
Batch size 160	97.201	60.57	51.011	44.23	44.131	44.265	44.2	44.229	43.615	43.979	43.655
Fetch time 16											
Batch size 4	105.481	66.119	53.034	46.815	45.171	44.013	43.623	43.623	43.62	43.62	43.615
Batch size 8	103.729	65.143	51.537	45.8	44.713	43.652	43.613	43.623	43.615	43.617	43.615
Batch size 16	101.83	63.657	50.158	45.06	44.049	43.609	43.614	43.615	43.615	43.615	43.615
Batch size 40	99.838	60.979	48.998	43.969	43.77	43.614	43.615	43.615	43.615	43.615	43.615
Batch size 80	96.914	58.512	47.502	43.282	43.394	43.665	43.649	43.735	43.615	43.794	43.665
Batch size 160	94.952	58.988	48.418	43.967	44.053	44.264	44.2	44.229	43.615	43.979	43.655
Fetch time 32											
Batch size 4	106.155	66.602	53.799	47.221	44.527	42.727	42.336	42.609	43.235	43.593	43.614
Batch size 8	104.101	65.488	52.235	46.245	44.08	42.394	42.358	42.623	43.252	43.605	43.614
Batch size 16	102.238	63.442	50.491	45.059	43.337	42.334	42.416	42.639	43.29	43.612	43.614
Batch size 40	98.747	61.039	48.803	44.041	42.977	42.54	42.66	42.794	43.391	43.615	43.614
Batch size 80	96.797	58.354	47.504	43.707	42.701	42.907	43.076	43.344	43.565	43.794	43.651
Batch size 160	95.56	58.234	48.023	43.924	43.899	44.209	44.195	44.209	43.615	44.024	43.655
Fetch time 64											
Batch size 4	106.528	67.17	53.854	47.269	44.426	42.57	42.142	42.17	42.224	42.289	42.61
Batch size 8	104.416	66.029	52.415	46.318	43.831	42.257	42.159	42.198	42.262	42.339	42.637
Batch size 16	102.302	63.946	50.393	45.07	43.04	42.158	42.222	42.272	42.36	42.479	42.708
Batch size 40	99.538	60.855	48.945	43.91	42.873	42.303	42.509	42.542	42.717	42.911	43.22
Batch size 80	96.534	58.395	47.984	43.436	42.532	42.802	43.041	43.248	43.453	43.732	43.611
Batch size 160	94.083	58.341	48.649	44.205	43.903	44.172	44.197	44.102	43.614	44.067	43.614
Fetch time 128											
Batch size 4	106.633	67.215	54.05	47.374	44.304	42.508	42.055	42.08	42.096	42.133	42.205
Batch size 8	104.446	66.059	52.287	46.112	43.782	42.211	42.078	42.114	42.136	42.186	42.28
Batch size 16	102.287	63.773	50.739	45.159	43.046	42.118	42.134	42.181	42.246	42.319	42.452
Batch size 40	99.762	60.971	48.726	43.902	42.849	42.377	42.315	42.416	42.577	42.759	43.17
Batch size 80	96.544	58.679	47.561	43.747	42.526	42.81	42.818	43.123	43.407	43.734	43.612
Batch size 160	94.895	60.33	48.078	44.09	44.062	44.092	44.094	44.014	43.611	44.067	43.612

Table 51: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the glimpse trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 4											
Batch size 4	26.128	18.833	17.411	15.52	13.862	12.667	11.554	10.437	10.249	10.26	10.252
Batch size 8	26.651	18.755	16.815	15.207	13.373	11.861	10.74	10.252	10.244	10.251	10.25
Batch size 16	26.177	17.831	16.035	14.473	12.608	11.452	10.244	10.244	10.244	10.244	10.246
Batch size 40	25.345	16.858	15.096	13.367	11.83	10.809	10.353	10.318	10.244	10.338	10.278
Batch size 80	24.775	16.712	14.936	13.224	11.23	10.582	10.591	10.577	10.448	10.597	10.569
Batch size 160	24.347	16.47	14.807	12.826	11.44	11.23	11.393	11.411	11.29	11.006	10.571
Fetch time 8											
Batch size 4	26.157	17.987	14.787	13.436	12.555	12.183	11.417	10.437	10.249	10.26	10.252
Batch size 8	26.642	17.867	14.172	12.895	12.013	11.394	10.74	10.252	10.244	10.251	10.25
Batch size 16	26.027	16.935	13.557	12.131	11.382	10.928	10.244	10.244	10.244	10.244	10.246
Batch size 40	25.348	16.289	13.392	11.525	10.767	10.666	10.353	10.318	10.244	10.338	10.278
Batch size 80	24.95	16.012	13.046	11.622	10.91	10.654	10.646	10.577	10.448	10.597	10.569
Batch size 160	24.377	16.05	13.599	12.282	11.615	11.358	11.393	11.411	11.29	11.006	10.571
Fetch time 16											
Batch size 4	26.082	18.02	14.686	12.761	11.182	10.599	10.147	9.886	9.959	10.138	10.252
Batch size 8	26.591	17.75	14.203	12.176	11.07	10.516	10.017	9.868	9.971	10.151	10.25
Batch size 16	25.982	17.166	13.636	11.806	10.683	10.341	9.927	9.854	10.01	10.18	10.246
Batch size 40	25.223	16.171	13.331	11.671	10.624	10.42	10.15	9.978	10.124	10.335	10.278
Batch size 80	24.825	15.921	13.057	11.805	11	10.664	10.587	10.61	10.46	10.597	10.569
Batch size 160	24.462	15.972	13.897	12.294	12.024	11.948	11.858	11.301	11.59	10.991	10.586
Fetch time 32											
Batch size 4	26.029	17.934	14.696	12.701	11.163	10.571	10.128	9.861	9.676	9.683	9.793
Batch size 8	26.351	17.837	14.292	12.17	10.954	10.434	10.118	9.816	9.678	9.699	9.817
Batch size 16	25.962	16.986	13.582	11.772	10.711	10.344	10.052	9.835	9.768	9.754	9.879
Batch size 40	25.27	16.282	13.347	11.847	10.747	10.447	10.341	10.008	9.995	10.052	10.129
Batch size 80	24.9	16.073	13.288	12.003	11.215	10.837	10.781	10.647	10.382	10.572	10.577
Batch size 160	24.435	16.01	13.66	12.866	12.324	12.356	12.229	11.772	11.535	11.097	10.826
Fetch time 64											
Batch size 4	26.029	17.934	14.68	12.656	11.224	10.504	10.166	9.834	9.692	9.767	9.677
Batch size 8	26.351	17.762	14.23	12.215	11.08	10.474	10.098	9.896	9.709	9.737	9.713
Batch size 16	26.244	17.018	13.503	11.91	10.825	10.375	10.111	10.009	9.768	9.746	9.794
Batch size 40	25.167	16.238	13.364	11.949	10.682	10.548	10.358	10.204	9.931	10.034	10.106
Batch size 80	24.888	16.188	13.216	12.046	11.189	10.97	10.776	10.645	10.605	10.708	10.67
Batch size 160	24.392	15.953	13.671	12.775	12.323	12.384	12.119	12.194	11.985	11.542	11.076
Fetch time 128											
Batch size 4	26.029	17.934	14.68	12.656	11.285	10.481	10.165	9.903	9.677	9.767	9.678
Batch size 8	26.351	17.762	14.278	12.19	11.05	10.442	10.129	9.862	9.701	9.735	9.714
Batch size 16	26.244	16.973	13.547	11.743	10.779	10.301	10.188	9.871	9.774	9.759	9.796
Batch size 40	25.219	16.114	13.346	11.759	10.749	10.471	10.515	10.203	10.178	10.06	10.088
Batch size 80	24.771	16.056	12.999	11.9	11.462	11.147	11.111	11.12	10.722	10.721	10.447
Batch size 160	24.377	16.148	13.651	12.722	12.605	12.135	12.557	12.579	12.059	11.549	10.615

Table 52: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the ld trace.

Disks	1	2	3	4	5	6
Fetch time 4						
Batch size 4	92.487	81.692	81.991	81.993	81.992	81.99
Batch size 8	90.455	81.691	81.988	81.986	81.982	81.978
Batch size 16	88.578	81.69	81.976	81.977	81.977	81.977
Batch size 40	85.763	81.871	82.253	82.25	81.977	82.234
Batch size 80	85.166	82.485	82.732	82.68	82.106	82.544
Batch size 160	85.002	83.489	83.497	83.284	82.401	83.086
Fetch time 8						
Batch size 4	92.487	81.163	81.359	81.694	81.965	81.986
Batch size 8	90.455	81.165	81.358	81.693	81.974	81.976
Batch size 16	88.578	81.169	81.357	81.7	81.975	81.975
Batch size 40	85.763	81.364	81.665	82.02	81.976	82.233
Batch size 80	85.166	82.024	82.204	82.525	82.105	82.543
Batch size 160	84.986	83.106	83.232	83.242	82.4	83.085
Fetch time 16						
Batch size 4	92.487	81.164	81.166	81.173	81.224	81.365
Batch size 8	90.455	81.166	81.17	81.177	81.224	81.362
Batch size 16	88.578	81.168	81.177	81.192	81.238	81.384
Batch size 40	85.763	81.363	81.514	81.568	81.352	81.71
Batch size 80	85.165	81.998	82.126	82.186	81.68	82.257
Batch size 160	84.99	83.067	83.165	83.141	82.342	83.04
Fetch time 32						
Batch size 4	92.487	81.164	81.165	81.17	81.173	81.177
Batch size 8	90.455	81.164	81.167	81.176	81.18	81.187
Batch size 16	88.578	81.164	81.171	81.193	81.208	81.226
Batch size 40	85.765	81.36	81.497	81.565	81.312	81.633
Batch size 80	85.165	81.984	82.073	82.116	81.689	82.219
Batch size 160	84.984	83.02	83.063	83.104	82.347	82.961
Fetch time 64						
Batch size 4	92.487	81.164	81.164	81.169	81.17	81.176
Batch size 8	90.455	81.167	81.166	81.172	81.175	81.18
Batch size 16	88.578	81.169	81.168	81.184	81.199	81.213
Batch size 40	85.765	81.365	81.494	81.523	81.293	81.604
Batch size 80	85.165	81.976	82.051	82.086	81.627	82.155
Batch size 160	84.985	83.015	83.014	83.091	82.346	83.054
Fetch time 128						
Batch size 4	92.487	81.165	81.166	81.172	81.171	81.175
Batch size 8	90.455	81.165	81.17	81.177	81.177	81.178
Batch size 16	88.578	81.169	81.172	81.18	81.195	81.213
Batch size 40	85.765	81.352	81.498	81.518	81.289	81.577
Batch size 80	85.168	81.982	82.059	82.082	81.617	82.155
Batch size 160	84.988	83.026	83.075	83.043	82.347	83.032

Table 53: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the postgres-join trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 4											
Batch size 4	52.54	26.927	20.823	16.389	13.787	13.277	13.275	13.273	13.269	13.28	13.272
Batch size 8	50.627	26.314	20.25	15.712	13.294	13.265	13.264	13.272	13.264	13.271	13.27
Batch size 16	49.017	25.066	19.493	14.959	13.264	13.264	13.264	13.264	13.264	13.264	13.266
Batch size 40	46.106	23.816	18.577	13.859	13.272	13.412	13.508	13.476	13.264	13.341	13.287
Batch size 80	43.782	22.778	17.716	13.878	13.622	13.906	13.88	13.821	13.313	13.618	13.641
Batch size 160	41.995	22.471	16.889	14.707	14.281	14.368	14.496	14.338	13.832	13.806	13.656
Fetch time 8											
Batch size 4	52.54	26.925	18.686	14.807	13.157	13.248	13.27	13.271	13.268	13.28	13.271
Batch size 8	50.606	26.313	18.011	14.236	13.131	13.246	13.263	13.27	13.263	13.271	13.269
Batch size 16	48.987	25.066	17.278	13.882	13.146	13.259	13.263	13.262	13.263	13.264	13.265
Batch size 40	46.075	23.816	16.586	13.181	13.209	13.412	13.507	13.474	13.263	13.341	13.286
Batch size 80	43.782	22.724	16.002	13.749	13.611	13.907	13.879	13.819	13.312	13.618	13.641
Batch size 160	41.995	21.496	15.797	14.672	14.253	14.368	14.494	14.336	13.831	13.806	13.655
Fetch time 16											
Batch size 4	52.557	26.925	18.687	14.807	13.078	13.052	13.05	13.048	13.134	13.255	13.272
Batch size 8	50.627	26.314	18.011	14.236	13.044	13.04	13.039	13.047	13.147	13.26	13.27
Batch size 16	49.017	25.066	17.278	13.882	13.039	13.039	13.039	13.051	13.185	13.263	13.266
Batch size 40	46.106	23.816	16.586	13.164	13.047	13.191	13.326	13.343	13.254	13.34	13.287
Batch size 80	43.782	22.724	16.001	13.654	13.432	13.779	13.814	13.814	13.314	13.617	13.641
Batch size 160	41.995	21.496	16.398	14.681	14.293	14.345	14.494	14.301	13.827	13.805	13.641
Fetch time 32											
Batch size 4	52.535	26.921	18.681	14.801	13.074	13.05	13.049	13.048	13.044	13.055	13.054
Batch size 8	50.599	26.311	18.007	14.236	13.044	13.04	13.039	13.047	13.039	13.046	13.068
Batch size 16	48.98	25.066	17.278	13.882	13.039	13.039	13.039	13.039	13.039	13.039	13.118
Batch size 40	46.039	23.816	16.586	13.164	13.047	13.187	13.283	13.255	13.07	13.198	13.27
Batch size 80	43.711	22.724	16.001	13.654	13.406	13.746	13.791	13.754	13.309	13.615	13.617
Batch size 160	41.987	21.496	16.333	14.638	14.208	14.37	14.504	14.184	13.746	13.603	13.638
Fetch time 64											
Batch size 4	52.531	26.918	18.682	14.802	13.073	13.047	13.045	13.044	13.041	13.053	13.047
Batch size 8	50.595	26.309	18.008	14.232	13.038	13.035	13.038	13.047	13.039	13.046	13.045
Batch size 16	48.98	25.062	17.275	13.882	13.039	13.039	13.039	13.039	13.039	13.039	13.042
Batch size 40	46.039	23.812	16.586	13.164	13.047	13.187	13.284	13.251	13.064	13.199	13.275
Batch size 80	43.711	22.716	16.001	13.653	13.424	13.779	13.777	13.748	13.317	13.6	13.588
Batch size 160	41.987	21.497	15.924	14.636	14.295	14.343	14.504	14.22	13.713	13.873	13.621
Fetch time 128											
Batch size 4	52.529	26.914	18.676	14.798	13.07	13.045	13.045	13.043	13.04	13.053	13.047
Batch size 8	50.595	26.309	18.003	14.228	13.036	13.033	13.034	13.042	13.036	13.044	13.045
Batch size 16	48.98	25.062	17.27	13.874	13.032	13.037	13.035	13.039	13.039	13.039	13.042
Batch size 40	46.039	23.812	16.578	13.158	13.044	13.187	13.283	13.251	13.064	13.22	13.243
Batch size 80	43.711	22.716	15.993	13.649	13.397	13.733	13.801	13.774	13.266	13.609	13.622
Batch size 160	41.987	21.492	16.415	14.582	14.199	14.736	14.555	14.492	13.752	13.975	13.633

Table 54: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the postgres-select trace.

Disks	1	2	3	4	5	6
Fetch time 4						
Batch size 4	66.434	33.368	39.661	34.022	34.103	34.097
Batch size 8	65.548	33.295	38.575	34.089	34.162	34.155
Batch size 16	65.02	33.423	37.646	34.181	34.242	34.181
Batch size 40	64.155	33.608	37.314	34.331	34.419	34.302
Batch size 80	64.275	33.807	37.418	34.563	34.701	34.741
Batch size 160	65.044	34.446	37.36	35.585	35.9	35.358
Fetch time 8						
Batch size 4	66.426	33.35	34.334	33.23	33.538	33.767
Batch size 8	65.723	33.262	33.891	33.29	33.606	33.833
Batch size 16	64.971	33.302	33.285	33.38	33.704	33.876
Batch size 40	64.184	33.493	33.643	33.563	33.938	34.04
Batch size 80	64.185	33.728	33.934	33.934	34.315	34.533
Batch size 160	64.759	34.379	34.967	35.23	35.802	35.293
Fetch time 16						
Batch size 4	66.423	33.348	34.616	33.138	33.048	33.127
Batch size 8	65.689	33.2	34.235	33.175	33.119	33.202
Batch size 16	64.974	33.3	33.477	33.269	33.235	33.276
Batch size 40	64.171	33.495	33.473	33.479	33.528	33.546
Batch size 80	64.336	33.731	33.866	33.834	34.058	34.226
Batch size 160	64.519	34.361	34.822	35.164	35.741	35.163
Fetch time 32						
Batch size 4	66.423	33.348	34.534	33.125	33.044	33.105
Batch size 8	65.658	33.215	34.142	33.171	33.116	33.177
Batch size 16	64.967	33.293	33.57	33.257	33.227	33.251
Batch size 40	64.109	33.485	33.488	33.457	33.514	33.485
Batch size 80	64.092	33.792	33.859	33.844	33.966	34.173
Batch size 160	64.713	34.339	34.797	35.126	35.731	35.034
Fetch time 64						
Batch size 4	66.423	33.348	34.534	33.125	33.042	33.105
Batch size 8	65.658	33.215	34.151	33.171	33.115	33.17
Batch size 16	64.967	33.293	33.433	33.252	33.223	33.229
Batch size 40	64.18	33.48	33.468	33.431	33.471	33.478
Batch size 80	64.199	33.688	33.813	33.796	33.951	34.162
Batch size 160	64.789	34.278	34.853	35.168	35.739	35.193
Fetch time 128						
Batch size 4	66.423	33.348	34.534	33.125	33.042	33.105
Batch size 8	65.658	33.215	34.151	33.165	33.115	33.179
Batch size 16	64.967	33.293	33.433	33.259	33.225	33.239
Batch size 40	64.18	33.486	33.48	33.44	33.474	33.463
Batch size 80	64.26	33.692	33.826	33.778	33.915	34.159
Batch size 160	64.924	34.299	34.813	35.175	35.662	35.198

Table 55: Reverse aggressive elapsed time as a function of fetch time estimate and batch size on the xds trace.

G Performance data: varying *fixed horizon*'s horizon

This section contains the performance data for *fixed horizon* with varying values of the horizon.

Disks	1	2	3	4	5	6
Horizon 16						
fetches	4716	4716	4716	4716	4716	4716
driver time (sec)	2.358	2.358	2.358	2.358	2.358	2.358
stall time (sec)	0.023	0.023	0.023	0.023	0.023	0.023
elapsed time (sec)	105.919	105.919	105.919	105.919	105.919	105.919
average fetch time (msec)	3.153	3.171	3.196	3.234	3.245	3.293
average disk utilization	0.14	0.071	0.047	0.036	0.029	0.024
Horizon 32						
fetches	4716	4716	4716	4716	4716	4716
driver time (sec)	2.358	2.358	2.358	2.358	2.358	2.358
stall time (sec)	0.022	0.022	0.022	0.022	0.022	0.022
elapsed time (sec)	105.918	105.918	105.918	105.918	105.918	105.918
average fetch time (msec)	3.145	3.182	3.201	3.241	3.259	3.294
average disk utilization	0.14	0.071	0.048	0.036	0.029	0.024
Horizon 64						
fetches	4789	4789	4789	4789	4789	4789
driver time (sec)	2.3945	2.3945	2.3945	2.3945	2.3945	2.3945
stall time (sec)	0.026	0.008	0.008	0.008	0.008	0.008
elapsed time (sec)	105.959	105.941	105.941	105.941	105.941	105.941
average fetch time (msec)	3.155	3.19	3.232	3.269	3.29	3.328
average disk utilization	0.14	0.072	0.049	0.037	0.03	0.025
Horizon 128						
fetches	5182	5182	5182	5182	5182	5182
driver time (sec)	2.591	2.591	2.591	2.591	2.591	2.591
stall time (sec)	0.249	0	0	0	0	0
elapsed time (sec)	106.378	106.129	106.129	106.129	106.129	106.129
average fetch time (msec)	3.171	3.208	3.256	3.286	3.32	3.37
average disk utilization	0.15	0.078	0.053	0.04	0.032	0.027
Horizon 256						
fetches	6005	6005	6005	6005	6005	6005
driver time (sec)	3.0025	3.0025	3.0025	3.0025	3.0025	3.0025
stall time (sec)	0.664	0	0.025	0	0	0
elapsed time (sec)	107.205	106.541	106.566	106.541	106.541	106.541
average fetch time (msec)	3.183	3.217	3.266	3.292	3.33	3.365
average disk utilization	0.18	0.091	0.061	0.046	0.038	0.032
Horizon 512						
fetches	8812	8812	8812	8812	8812	8812
driver time (sec)	4.406	4.406	4.406	4.406	4.406	4.406
stall time (sec)	1.444	0	0.205	0.04	0	0
elapsed time (sec)	109.388	107.944	108.149	107.984	107.944	107.944
average fetch time (msec)	3.147	3.163	3.187	3.204	3.214	3.228
average disk utilization	0.25	0.13	0.087	0.065	0.052	0.044
Horizon 1024						
fetches	8812	8812	8812	8812	8812	8812
driver time (sec)	4.406	4.406	4.406	4.406	4.406	4.406
stall time (sec)	1.533	0	0.219	0.055	0	0
elapsed time (sec)	109.477	107.944	108.163	107.999	107.944	107.944
average fetch time (msec)	3.148	3.166	3.185	3.206	3.215	3.231
average disk utilization	0.25	0.13	0.086	0.065	0.052	0.044
Horizon 2048						
fetches	8812	8812	8812	8812	8812	8812
driver time (sec)	4.406	4.406	4.406	4.406	4.406	4.406
stall time (sec)	1.533	0	0.219	0.055	0	0
elapsed time (sec)	109.477	107.944	108.163	107.999	107.944	107.944
average fetch time (msec)	3.148	3.166	3.185	3.206	3.215	3.231
average disk utilization	0.25	0.13	0.086	0.065	0.052	0.044

Table 56: Fixed horizon performance as a function of horizon on the dinero trace.

Disks	1	2	3	4	5	6
Horizon 16						
fetches	4953	4953	4953	4953	4953	4953
driver time (sec)	2.4765	2.4765	2.4765	2.4765	2.4765	2.4765
stall time (sec)	3.476	0.022	0.022	0.022	0.022	0.022
elapsed time (sec)	30.887	27.433	27.433	27.433	27.433	27.433
average fetch time (msec)	3.543	3.249	3.255	3.257	3.295	3.311
average disk utilization	0.57	0.29	0.2	0.15	0.12	0.1
Horizon 32						
fetches	4953	4953	4953	4953	4953	4953
driver time (sec)	2.4765	2.4765	2.4765	2.4765	2.4765	2.4765
stall time (sec)	3.38	0.022	0.022	0.022	0.022	0.022
elapsed time (sec)	30.791	27.433	27.433	27.433	27.433	27.433
average fetch time (msec)	3.51	3.237	3.242	3.272	3.283	3.326
average disk utilization	0.56	0.29	0.2	0.15	0.12	0.1
Horizon 64						
fetches	4959	4959	4959	4959	4959	4959
driver time (sec)	2.4795	2.4795	2.4795	2.4795	2.4795	2.4795
stall time (sec)	3.121	0.012	0.012	0.012	0.012	0.012
elapsed time (sec)	30.535	27.426	27.426	27.426	27.426	27.426
average fetch time (msec)	3.524	3.251	3.277	3.301	3.333	3.368
average disk utilization	0.57	0.29	0.2	0.15	0.12	0.1
Horizon 128						
fetches	5471	5471	5471	5471	5471	5471
driver time (sec)	2.7355	2.7355	2.7355	2.7355	2.7355	2.7355
stall time (sec)	3.107	0	0	0	0	0
elapsed time (sec)	30.777	27.67	27.67	27.67	27.67	27.67
average fetch time (msec)	3.558	3.373	3.405	3.404	3.439	3.472
average disk utilization	0.63	0.33	0.22	0.17	0.14	0.11
Horizon 256						
fetches	6059	6059	6059	6059	6059	6059
driver time (sec)	3.0295	3.0295	3.0295	3.0295	3.0295	3.0295
stall time (sec)	2.726	0.135	0	0	0	0
elapsed time (sec)	30.69	28.099	27.964	27.964	27.964	27.964
average fetch time (msec)	3.624	3.398	3.427	3.424	3.464	3.504
average disk utilization	0.72	0.37	0.25	0.19	0.15	0.13
Horizon 512						
fetches	8535	8535	8535	8535	8535	8535
driver time (sec)	4.2675	4.2675	4.2675	4.2675	4.2675	4.2675
stall time (sec)	5.01	0.487	0.198	0	0	0
elapsed time (sec)	34.212	29.689	29.4	29.202	29.202	29.202
average fetch time (msec)	3.751	3.318	3.354	3.318	3.363	3.38
average disk utilization	0.94	0.48	0.32	0.24	0.2	0.16
Horizon 1024						
fetches	8535	8535	8535	8535	8535	8535
driver time (sec)	4.2675	4.2675	4.2675	4.2675	4.2675	4.2675
stall time (sec)	5.126	0.521	0.216	0	0	0
elapsed time (sec)	34.328	29.723	29.418	29.202	29.202	29.202
average fetch time (msec)	3.759	3.349	3.37	3.327	3.369	3.383
average disk utilization	0.93	0.48	0.33	0.24	0.2	0.16
Horizon 2048						
fetches	8535	8535	8535	8535	8535	8535
driver time (sec)	4.2675	4.2675	4.2675	4.2675	4.2675	4.2675
stall time (sec)	5.126	0.521	0.216	0	0	0
elapsed time (sec)	34.328	29.723	29.418	29.202	29.202	29.202
average fetch time (msec)	3.759	3.349	3.37	3.327	3.369	3.383
average disk utilization	0.93	0.48	0.33	0.24	0.2	0.16

Table 57: Fixed horizon performance as a function of horizon on the cscope1 trace.

Disks	1	2	3	4	5	6
Horizon 16						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	37.752	28.281	21.894	16.112	13.633	11.645
elapsed time (sec)	77.844	68.373	61.986	56.204	53.725	51.737
average fetch time (msec)	10.09	16.023	17.859	18.343	18.647	19.025
average disk utilization	0.77	0.7	0.57	0.49	0.41	0.37
Horizon 32						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	35.359	25.027	17.587	12.245	8.758	6.44
elapsed time (sec)	75.451	65.119	57.679	52.337	48.85	46.532
average fetch time (msec)	9.793	15.518	17.596	18.151	18.57	18.739
average disk utilization	0.77	0.71	0.61	0.52	0.45	0.4
Horizon 64						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	32.366	21.956	14.063	8.965	5.545	3.756
elapsed time (sec)	72.458	62.048	54.155	49.057	45.637	43.848
average fetch time (msec)	9.393	14.97	17.165	18.012	18.49	18.971
average disk utilization	0.77	0.72	0.63	0.55	0.48	0.43
Horizon 128						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	30.072	19.129	11.066	6.34	3.078	1.395
elapsed time (sec)	70.164	59.221	51.158	46.432	43.17	41.487
average fetch time (msec)	9.201	14.447	16.555	17.684	18.147	18.608
average disk utilization	0.78	0.73	0.64	0.57	0.5	0.45
Horizon 256						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	25.667	14.606	7.443	4.047	1.299	0.076
elapsed time (sec)	65.759	54.698	47.535	44.139	41.391	40.168
average fetch time (msec)	8.94	13.632	15.803	17.261	17.691	18.401
average disk utilization	0.81	0.74	0.66	0.58	0.51	0.46
Horizon 512						
fetches	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	19.192	8.934	3.017	1.164	0.208	0.169
elapsed time (sec)	59.284	49.026	43.109	41.256	40.3	40.261
average fetch time (msec)	8.456	12.84	15.156	16.622	17.596	18.407
average disk utilization	0.85	0.78	0.7	0.6	0.52	0.45
Horizon 1024						
fetches	6736	6736	6736	6736	6736	6736
driver time (sec)	3.368	3.368	3.368	3.368	3.368	3.368
stall time (sec)	19.012	7.313	2.32	0.787	0.623	0.484
elapsed time (sec)	59.489	47.79	42.797	41.264	41.1	40.961
average fetch time (msec)	8.158	12.231	14.679	16.219	17.298	18.055
average disk utilization	0.92	0.86	0.77	0.66	0.57	0.49
Horizon 2048						
fetches	8299	8299	8299	8299	8299	8299
driver time (sec)	4.1495	4.1495	4.1495	4.1495	4.1495	4.1495
stall time (sec)	23.94	8.745	2.803	2.103	1.709	1.176
elapsed time (sec)	65.199	50.004	44.062	43.362	42.968	42.435
average fetch time (msec)	7.744	11.784	14.304	16.248	17.258	18.029
average disk utilization	0.99	0.98	0.9	0.78	0.67	0.59

Table 58: Fixed horizon performance as a function of horizon on the cscope2 trace.

Disks	1	2	3	4	5	6
Horizon 16						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	36.253	16.646	9.658	7.205	5.024	3.906
elapsed time (sec)	49.273	29.666	22.678	20.225	18.044	16.926
average fetch time (msec)	15.493	15.727	15.617	15.76	15.508	15.601
average disk utilization	0.97	0.82	0.71	0.6	0.53	0.47
Horizon 32						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	34.453	14.462	7.778	4.953	2.902	1.863
elapsed time (sec)	47.473	27.482	20.798	17.973	15.922	14.883
average fetch time (msec)	14.914	15.289	15.441	15.526	15.415	15.429
average disk utilization	0.97	0.86	0.76	0.67	0.6	0.53
Horizon 64						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	32.46	12.642	5.718	3.059	1.206	0.51
elapsed time (sec)	45.48	25.662	18.738	16.079	14.226	13.53
average fetch time (msec)	14.407	14.853	15.033	15.11	15.217	15.386
average disk utilization	0.98	0.89	0.83	0.72	0.66	0.58
Horizon 128						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.405	10.83	4.525	1.745	0.422	0.064
elapsed time (sec)	43.425	23.85	17.545	14.765	13.442	13.084
average fetch time (msec)	13.831	14.252	14.587	14.67	14.925	15.255
average disk utilization	0.98	0.92	0.85	0.77	0.69	0.6
Horizon 256						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	28.591	9.283	3.336	1.104	0.183	0
elapsed time (sec)	41.611	22.303	16.356	14.124	13.203	13.02
average fetch time (msec)	13.281	13.506	13.678	14.291	14.633	15.212
average disk utilization	0.98	0.93	0.86	0.78	0.68	0.6
Horizon 512						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	27.117	8.234	3.284	1.14	0.382	0.53
elapsed time (sec)	40.137	21.254	16.304	14.16	13.402	13.55
average fetch time (msec)	12.586	12.687	13.286	13.586	14.463	15.001
average disk utilization	0.97	0.92	0.84	0.74	0.67	0.57
Horizon 1024						
fetches	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	23.842	6.768	2.743	1.827	1.116	1.066
elapsed time (sec)	36.862	19.788	15.763	14.847	14.136	14.086
average fetch time (msec)	11.305	11.583	12.206	13.098	13.752	13.99
average disk utilization	0.95	0.9	0.8	0.68	0.6	0.51
Horizon 2048						
fetches	3572	3572	3572	3572	3572	3572
driver time (sec)	1.786	1.786	1.786	1.786	1.786	1.786
stall time (sec)	24.368	7.001	3.405	2.278	1.298	1.171
elapsed time (sec)	37.632	20.265	16.669	15.542	14.562	14.435
average fetch time (msec)	10.116	10.384	11.257	11.976	12.788	13.109
average disk utilization	0.96	0.92	0.8	0.69	0.63	0.54

Table 59: Fixed horizon performance as a function of horizon on the postgres-select trace.

H Performance data: *forestall* with a fixed fetch time estimate

This section contains the performance data for *forestall* with a static fetch time estimate.

Disks	1	2	3	4	5	6
Fetch time 2						
fetches	8812	4753	4753	4753	4753	4753
driver time (sec)	4.406	2.3765	2.3765	2.3765	2.3765	2.3765
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	105.915	105.915	105.915	105.915	105.916
average fetch time (msec)	3.173	3.208	3.254	3.283	3.297	3.324
average disk utilization	0.26	0.072	0.049	0.037	0.03	0.025
Fetch time 4						
fetches	8812	10268	4909	4753	4753	4753
driver time (sec)	4.406	5.134	2.4545	2.3765	2.3765	2.3765
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	108.672	105.993	105.915	105.915	105.916
average fetch time (msec)	3.143	3.857	3.341	3.283	3.297	3.324
average disk utilization	0.26	0.18	0.052	0.037	0.03	0.025
Fetch time 8						
fetches	8812	8818	8838	10277	4948	4852
driver time (sec)	4.406	4.409	4.419	5.1385	2.474	2.426
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	107.947	107.957	108.677	106.012	105.965
average fetch time (msec)	3.143	3.152	3.185	3.943	3.551	3.343
average disk utilization	0.26	0.13	0.087	0.093	0.033	0.026
Fetch time 15						
fetches	8812	8815	8844	8821	8830	8824
driver time (sec)	4.406	4.4075	4.422	4.4105	4.415	4.412
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	107.946	107.96	107.949	107.953	107.951
average fetch time (msec)	3.141	3.149	3.182	3.182	3.194	3.206
average disk utilization	0.26	0.13	0.087	0.065	0.052	0.044
Fetch time 30						
fetches	8812	8812	8832	8824	8816	8821
driver time (sec)	4.406	4.406	4.416	4.412	4.408	4.4105
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	107.944	107.954	107.95	107.946	107.95
average fetch time (msec)	3.142	3.147	3.177	3.182	3.19	3.204
average disk utilization	0.26	0.13	0.087	0.065	0.052	0.044
Fetch time 60						
fetches	8812	8812	8823	8816	8819	8825
driver time (sec)	4.406	4.406	4.4115	4.408	4.4095	4.4125
stall time (sec)	0.145	0	0	0	0	0.001
elapsed time (sec)	108.089	107.944	107.95	107.946	107.948	107.952
average fetch time (msec)	3.141	3.146	3.176	3.177	3.189	3.207
average disk utilization	0.26	0.13	0.087	0.065	0.052	0.044

Table 60: Forestall performance as a function of static fetch time estimate on the dinero trace.

Disks	1	2	3	4	5	6
Fetch time 2						
fetches	6892	4953	4953	4953	4953	4953
driver time (sec)	3.446	2.4765	2.4765	2.4765	2.4765	2.4765
stall time (sec)	0.782	0	0	0	0	0.001
elapsed time (sec)	29.162	27.411	27.411	27.411	27.411	27.412
average fetch time (msec)	3.74	3.243	3.321	3.295	3.331	3.342
average disk utilization	0.88	0.29	0.2	0.15	0.12	0.1
Fetch time 4						
fetches	6931	8656	5108	4953	4953	4953
driver time (sec)	3.4655	4.328	2.554	2.4765	2.4765	2.4765
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.262	27.488	27.411	27.411	27.412
average fetch time (msec)	3.753	3.57	3.507	3.295	3.331	3.342
average disk utilization	0.89	0.53	0.22	0.15	0.12	0.1
Fetch time 8						
fetches	6931	8570	8680	9650	5181	5063
driver time (sec)	3.4655	4.285	4.34	4.825	2.5905	2.5315
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.219	29.274	29.759	27.525	27.467
average fetch time (msec)	3.758	3.36	3.448	3.976	3.57	3.449
average disk utilization	0.89	0.49	0.34	0.32	0.13	0.11
Fetch time 15						
fetches	6931	8570	8676	8680	8623	8582
driver time (sec)	3.4655	4.285	4.338	4.34	4.3115	4.291
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.219	29.272	29.274	29.246	29.226
average fetch time (msec)	3.759	3.362	3.438	3.368	3.394	3.359
average disk utilization	0.89	0.49	0.34	0.25	0.2	0.16
Fetch time 30						
fetches	6931	8571	8673	8688	8623	8577
driver time (sec)	3.4655	4.2855	4.3365	4.344	4.3115	4.2885
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.22	29.271	29.278	29.246	29.224
average fetch time (msec)	3.759	3.365	3.427	3.366	3.389	3.358
average disk utilization	0.89	0.49	0.34	0.25	0.2	0.16
Fetch time 60						
fetches	6931	8570	8673	8681	8627	8583
driver time (sec)	3.4655	4.285	4.3365	4.3405	4.3135	4.2915
stall time (sec)	0.911	0	0	0	0	0.001
elapsed time (sec)	29.311	29.219	29.271	29.275	29.248	29.227
average fetch time (msec)	3.758	3.362	3.43	3.363	3.394	3.36
average disk utilization	0.89	0.49	0.34	0.25	0.2	0.16

Table 61: Forestall performance as a function of static fetch time estimate on the cscopel trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 2											
fetches	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966
driver time (sec)	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	33.592	23.879	15.09	8.553	5.27	3.226	1.792	1.258	0.812	0.262	0.018
elapsed time (sec)	73.684	63.971	55.182	48.645	45.362	43.318	41.884	41.35	40.904	40.354	40.11
average fetch time (msec)	9.476	15.459	17.353	18.156	18.318	18.707	18.866	19.181	19.048	19.212	19.31
average disk utilization	0.77	0.72	0.63	0.56	0.48	0.43	0.38	0.35	0.28	0.24	0.18
Fetch time 4											
fetches	6166	5966	5966	5966	5966	5966	5966	5966	5966	5966	5966
driver time (sec)	3.083	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	20.969	24.945	15.09	9.212	5.27	3.226	1.792	1.258	0.812	0.262	0.018
elapsed time (sec)	61.161	65.037	55.182	49.304	45.362	43.318	41.884	41.35	40.904	40.354	40.11
average fetch time (msec)	8.827	15.325	17.353	18.085	18.318	18.707	18.866	19.181	19.048	19.212	19.31
average disk utilization	0.89	0.7	0.63	0.55	0.48	0.43	0.38	0.35	0.28	0.24	0.18
Fetch time 8											
fetches	6284	6144	6025	5967	5966	5966	5966	5966	5966	5966	5966
driver time (sec)	3.142	3.072	3.0125	2.9835	2.983	2.983	2.983	2.983	2.983	2.983	2.983
stall time (sec)	15.971	14.011	20.831	11.275	7.757	3.371	1.912	1.258	0.812	0.262	0.018
elapsed time (sec)	56.222	54.192	60.953	51.368	47.849	43.463	42.004	41.35	40.904	40.354	40.11
average fetch time (msec)	8.768	13.401	14.909	17.84	18.302	18.699	18.803	19.191	19.048	19.212	19.31
average disk utilization	0.98	0.76	0.49	0.52	0.46	0.43	0.38	0.35	0.28	0.24	0.18
Fetch time 15											
fetches	6318	6333	6613	6036	5990	5969	5966	5966	5966	5966	5966
driver time (sec)	3.159	3.1665	3.3065	3.018	2.995	2.9845	2.983	2.983	2.983	2.983	2.983
stall time (sec)	15.858	5.998	2.274	6.358	5.106	1.933	2.128	0.746	0.438	0.321	0.018
elapsed time (sec)	56.126	46.274	42.69	46.485	45.21	42.027	42.22	40.838	40.53	40.413	40.11
average fetch time (msec)	8.773	13.294	14.556	16.693	17.083	18.439	18.61	19.06	19.131	19.294	19.304
average disk utilization	0.99	0.91	0.75	0.54	0.45	0.44	0.38	0.35	0.28	0.24	0.18
Fetch time 30											
fetches	6318	6592	7372	7298	7256	6664	6253	5997	5969	5970	5970
driver time (sec)	3.159	3.296	3.686	3.649	3.628	3.332	3.1265	2.9985	2.9845	2.985	2.985
stall time (sec)	15.858	5.597	1.798	0	0	0.002	0.16	0.009	0.023	0.02	0.025
elapsed time (sec)	56.126	46.002	42.593	40.758	40.737	40.443	40.396	40.117	40.117	40.114	40.119
average fetch time (msec)	8.773	13.256	14.494	16.687	16.764	17.98	18.135	18.889	19.126	19.144	19.305
average disk utilization	0.99	0.95	0.84	0.75	0.6	0.49	0.4	0.35	0.28	0.24	0.18
Fetch time 60											
fetches	6318	6592	7683	7857	8513	8152	7922	7607	7136	6672	6157
driver time (sec)	3.159	3.296	3.8415	3.9285	4.2565	4.076	3.961	3.8035	3.568	3.336	3.0785
stall time (sec)	15.858	5.597	1.798	0	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	56.126	46.002	42.749	41.038	41.366	41.186	41.07	40.922	40.682	40.446	40.188
average fetch time (msec)	8.773	13.257	14.436	16.575	16.88	17.901	18.191	18.771	19.019	19.104	19.226
average disk utilization	0.99	0.95	0.86	0.79	0.69	0.59	0.5	0.44	0.33	0.26	0.18

Table 62: Fore stall performance as a function of static fetch time estimate on the cscope2 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 2											
fetches	11877	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739
driver time (sec)	5.9385	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	32.201	14.155	6.993	2.386	1.123	0.395	0.18	0.067	0.085	0.001	0
elapsed time (sec)	112.24	94.125	86.963	82.356	81.093	80.365	80.15	80.037	80.055	79.971	79.97
average fetch time (msec)	7.782	12.15	14.801	16.136	16.856	17.427	17.847	18.21	18.622	18.753	19.17
average disk utilization	0.82	0.76	0.67	0.58	0.49	0.42	0.37	0.33	0.27	0.23	0.18
Fetch time 4											
fetches	11989	11739	11739	11739	11739	11739	11739	11739	11739	11739	11739
driver time (sec)	5.9945	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	15.255	19.98	7.427	2.386	1.123	0.395	0.18	0.067	0.085	0.001	0
elapsed time (sec)	95.35	99.95	87.397	82.356	81.093	80.365	80.15	80.037	80.055	79.971	79.97
average fetch time (msec)	7.687	11.882	14.786	16.136	16.856	17.427	17.847	18.21	18.622	18.753	19.17
average disk utilization	0.97	0.7	0.66	0.58	0.49	0.42	0.37	0.33	0.27	0.23	0.18
Fetch time 8											
fetches	12029	12380	11935	11739	11739	11739	11739	11739	11739	11739	11739
driver time (sec)	6.0145	6.19	5.9675	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	14.198	3.272	15.339	5.365	2.089	0.439	0.18	0.067	0.085	0.001	0
elapsed time (sec)	94.313	83.563	95.407	85.335	82.059	80.409	80.15	80.037	80.055	79.971	79.97
average fetch time (msec)	7.697	11.635	13.861	16.069	16.776	17.444	17.847	18.21	18.622	18.753	19.17
average disk utilization	0.98	0.86	0.58	0.55	0.48	0.42	0.37	0.33	0.27	0.23	0.18
Fetch time 15											
fetches	12069	13014	13732	12759	12118	11739	11739	11739	11739	11739	11739
driver time (sec)	6.0345	6.507	6.866	6.3795	6.059	5.8695	5.8695	5.8695	5.8695	5.8695	5.8695
stall time (sec)	13.943	2.862	0.64	0.052	0.935	0.188	0.335	0.009	0.084	0.001	0
elapsed time (sec)	94.078	83.47	81.607	80.532	81.095	80.158	80.305	79.979	80.054	79.971	79.97
average fetch time (msec)	7.703	11.602	13.64	15.752	16.266	17.462	17.749	18.334	18.598	18.749	19.162
average disk utilization	0.99	0.9	0.77	0.62	0.49	0.43	0.37	0.34	0.27	0.23	0.18
Fetch time 30											
fetches	12092	13414	14940	14520	14134	13588	13184	12677	12078	11749	11742
driver time (sec)	6.046	6.707	7.47	7.26	7.067	6.794	6.592	6.3385	6.039	5.8745	5.871
stall time (sec)	13.943	2.862	0.64	0.052	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	94.09	83.67	82.211	81.413	81.168	80.896	80.693	80.448	80.145	79.976	79.972
average fetch time (msec)	7.741	11.586	13.625	15.76	16.23	17.39	17.76	18.613	18.804	18.924	19.105
average disk utilization	0.99	0.93	0.83	0.7	0.57	0.49	0.41	0.37	0.28	0.23	0.18
Fetch time 60											
fetches	12092	13534	15442	15702	15780	15120	14760	14393	13977	13574	12798
driver time (sec)	6.046	6.767	7.721	7.851	7.89	7.56	7.38	7.1965	6.9885	6.787	6.399
stall time (sec)	13.943	2.862	0.64	0.052	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	94.09	83.73	82.462	82.004	81.991	81.662	81.481	81.306	81.094	80.889	80.5
average fetch time (msec)	7.741	11.585	13.8	15.782	16.244	17.301	17.723	18.584	18.766	18.925	19.119
average disk utilization	0.99	0.94	0.86	0.76	0.63	0.53	0.46	0.41	0.32	0.26	0.19

Table 63: Forestall performance as a function of static fetch time estimate on the cscope3 trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 2											
fetches	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493	6493
driver time (sec)	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465	3.2465
stall time (sec)	64.134	29.651	18.465	11.895	8.142	5.704	3.725	3.291	1.773	0.821	0.507
elapsed time (sec)	106.097	71.614	60.428	53.858	50.105	47.667	45.688	45.254	43.736	42.784	42.47
average fetch time (msec)	13.314	15.231	16.228	17.452	17.993	18.334	18.457	18.528	18.598	18.642	18.707
average disk utilization	0.81	0.69	0.58	0.53	0.47	0.42	0.37	0.33	0.28	0.24	0.18
Fetch time 4											
fetches	6531	6495	6521	6493	6493	6500	6493	6493	6493	6493	6493
driver time (sec)	3.2655	3.2475	3.2605	3.2465	3.2465	3.25	3.2465	3.2465	3.2465	3.2465	3.2465
stall time (sec)	57.697	31.749	19.621	11.895	8.142	5.821	3.725	3.306	1.773	0.821	0.507
elapsed time (sec)	99.679	73.713	61.598	53.858	50.105	47.787	45.688	45.269	43.736	42.784	42.47
average fetch time (msec)	13.103	15.003	16.19	17.452	17.993	18.33	18.457	18.524	18.598	18.626	18.707
average disk utilization	0.86	0.66	0.57	0.53	0.47	0.42	0.37	0.33	0.28	0.24	0.18
Fetch time 8											
fetches	6531	6578	6538	6503	6493	6497	6493	6493	6493	6493	6493
driver time (sec)	3.2655	3.289	3.269	3.2515	3.2465	3.2485	3.2465	3.2465	3.2465	3.2465	3.2465
stall time (sec)	56.363	24.886	35.324	16.784	11.185	6.451	4.07	3.321	1.818	0.821	0.507
elapsed time (sec)	98.345	66.891	77.309	58.752	53.148	48.416	46.033	45.284	43.781	42.784	42.47
average fetch time (msec)	13.104	14.396	15.337	17.336	17.923	18.361	18.455	18.533	18.614	18.628	18.707
average disk utilization	0.87	0.71	0.43	0.48	0.44	0.41	0.37	0.33	0.28	0.24	0.18
Fetch time 15											
fetches	6531	6647	6688	6538	6530	6505	6494	6493	6493	6493	6493
driver time (sec)	3.2655	3.3235	3.344	3.269	3.265	3.2525	3.247	3.2465	3.2465	3.2465	3.2465
stall time (sec)	55.448	20.658	12.116	7.584	12.461	6.181	5.497	2.676	1.863	1.076	0.617
elapsed time (sec)	97.43	62.698	54.176	49.569	54.442	48.15	47.46	44.639	43.826	43.039	42.58
average fetch time (msec)	13.104	14.423	14.918	16.525	17.05	18.137	18.457	18.516	18.585	18.648	18.737
average disk utilization	0.88	0.76	0.61	0.54	0.41	0.41	0.36	0.34	0.28	0.23	0.18
Fetch time 30											
fetches	6565	6687	6891	6769	6723	6616	6560	6514	6493	6493	6493
driver time (sec)	3.2825	3.3435	3.4455	3.3845	3.3615	3.308	3.28	3.257	3.2465	3.2465	3.2465
stall time (sec)	55.586	19.562	6.609	2.522	1.939	0.671	0.813	0.072	0.027	0.119	0.215
elapsed time (sec)	97.585	61.622	48.771	44.623	44.017	42.695	42.809	42.045	41.99	42.082	42.178
average fetch time (msec)	13.059	14.378	14.694	16.438	16.683	17.776	17.929	18.496	18.549	18.526	18.672
average disk utilization	0.88	0.78	0.69	0.62	0.51	0.46	0.39	0.36	0.29	0.24	0.18
Fetch time 60											
fetches	6610	6687	7087	6911	6969	6823	6964	6836	6703	6624	6565
driver time (sec)	3.305	3.3435	3.5435	3.4555	3.4845	3.4115	3.482	3.418	3.3515	3.312	3.2825
stall time (sec)	54.845	19.089	5.792	2.521	1.062	0.099	0.014	0.009	0.005	0.001	0
elapsed time (sec)	96.866	61.149	48.052	44.693	43.263	42.227	42.212	42.143	42.073	42.029	41.999
average fetch time (msec)	12.998	14.374	14.565	16.422	16.75	17.762	17.634	18.282	18.499	18.486	18.626
average disk utilization	0.89	0.79	0.72	0.63	0.54	0.48	0.42	0.37	0.29	0.24	0.18

Table 64: Forestall performance as a function of static fetch time estimate on the glimpse trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 2											
fetches	2900	2903	2903	2903	2903	2903	2903	2903	2903	2903	2903
driver time (sec)	1.45	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515
stall time (sec)	15.617	7.475	4.441	2.739	1.88	1.368	1.069	0.968	0.567	0.327	0.222
elapsed time (sec)	25.232	17.091	14.057	12.355	11.496	10.984	10.685	10.584	10.183	9.943	9.838
average fetch time (msec)	8.451	11.126	13.266	15.004	16.194	16.777	17.359	18.041	18.669	18.998	19.113
average disk utilization	0.97	0.94	0.91	0.88	0.82	0.74	0.67	0.62	0.53	0.46	0.35
Fetch time 4											
fetches	2981	2903	2903	2903	2903	2903	2903	2903	2903	2903	2903
driver time (sec)	1.4905	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515
stall time (sec)	15.245	8.724	4.8	2.835	1.88	1.368	1.069	0.968	0.567	0.327	0.222
elapsed time (sec)	24.9	18.34	14.416	12.451	11.496	10.984	10.685	10.584	10.183	9.943	9.838
average fetch time (msec)	8.248	10.959	13.249	15	16.194	16.777	17.359	18.041	18.669	18.998	19.113
average disk utilization	0.99	0.87	0.89	0.87	0.82	0.74	0.67	0.62	0.53	0.46	0.35
Fetch time 8											
fetches	2981	2982	3093	2903	2903	2903	2903	2903	2903	2903	2903
driver time (sec)	1.4905	1.491	1.5465	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515	1.4515
stall time (sec)	15.245	6.329	3.461	3.358	2.261	1.374	1.204	0.983	0.567	0.327	0.222
elapsed time (sec)	24.9	15.985	13.172	12.974	11.877	10.99	10.82	10.599	10.183	9.943	9.838
average fetch time (msec)	8.248	10.583	12.135	14.947	16.031	16.705	17.329	18.036	18.669	18.998	19.113
average disk utilization	0.99	0.99	0.95	0.84	0.78	0.74	0.66	0.62	0.53	0.46	0.35
Fetch time 15											
fetches	2981	2982	3137	3102	3302	2928	2909	2903	2903	2903	2903
driver time (sec)	1.4905	1.491	1.5685	1.551	1.651	1.464	1.4545	1.4515	1.4515	1.4515	1.4515
stall time (sec)	15.245	6.329	3.433	2.056	0.685	0.86	1.511	1.028	0.627	0.38	0.327
elapsed time (sec)	24.9	15.985	13.166	11.772	10.501	10.489	11.13	10.644	10.243	9.996	9.943
average fetch time (msec)	8.248	10.583	12.033	14.203	14.993	16.505	17.156	17.972	18.665	18.985	19.008
average disk utilization	0.99	0.99	0.96	0.94	0.94	0.77	0.64	0.61	0.53	0.46	0.35
Fetch time 30											
fetches	2981	2982	3137	3102	3310	3505	3737	3663	3448	3017	2917
driver time (sec)	1.4905	1.491	1.5685	1.551	1.655	1.7525	1.8685	1.8315	1.724	1.5085	1.4585
stall time (sec)	15.245	6.329	3.433	2.052	0.579	0.265	0.164	0.022	0.019	0.018	0.101
elapsed time (sec)	24.9	15.985	13.166	11.768	10.399	10.182	10.197	10.018	9.908	9.691	9.724
average fetch time (msec)	8.248	10.583	12.037	14.199	14.932	15.957	16.449	17.305	18.174	18.925	18.949
average disk utilization	0.99	0.99	0.96	0.94	0.95	0.92	0.86	0.79	0.63	0.49	0.36
Fetch time 60											
fetches	2981	2982	3137	3102	3310	3505	3734	3779	4080	4137	3880
driver time (sec)	1.4905	1.491	1.5685	1.551	1.655	1.7525	1.867	1.8895	2.04	2.0685	1.94
stall time (sec)	15.245	6.329	3.433	2.052	0.579	0.265	0.023	0.009	0.005	0.001	0
elapsed time (sec)	24.9	15.985	13.166	11.768	10.399	10.182	10.055	10.063	10.21	10.234	10.105
average fetch time (msec)	8.248	10.583	12.037	14.199	14.932	15.958	16.446	17.181	17.814	18.076	18.644
average disk utilization	0.99	0.99	0.96	0.94	0.95	0.92	0.87	0.81	0.71	0.61	0.45

Table 65: Fore stall performance as a function of static fetch time estimate on the ld trace.

Disks	1	2	3	4	5	6
Fetch time 2						
fetches	3855	3856	3855	3856	3855	3856
driver time (sec)	1.9275	1.928	1.9275	1.928	1.9275	1.928
stall time (sec)	4.765	0.152	0.258	0	0	0.001
elapsed time (sec)	85.908	81.296	81.401	81.144	81.143	81.145
average fetch time (msec)	16.863	17.094	17.051	17.666	17.507	17.701
average disk utilization	0.76	0.41	0.27	0.21	0.17	0.14
Fetch time 4						
fetches	4108	3856	3855	3856	3855	3856
driver time (sec)	2.054	1.928	1.9275	1.928	1.9275	1.928
stall time (sec)	3.994	0.152	0.258	0	0	0.001
elapsed time (sec)	85.264	81.296	81.401	81.144	81.143	81.145
average fetch time (msec)	14.79	16.841	17.051	17.666	17.507	17.701
average disk utilization	0.71	0.4	0.27	0.21	0.17	0.14
Fetch time 8						
fetches	4695	4174	3937	3856	3855	3856
driver time (sec)	2.3475	2.087	1.9685	1.928	1.9275	1.928
stall time (sec)	3.995	0.152	0.293	0	0	0.001
elapsed time (sec)	85.558	81.455	81.477	81.144	81.143	81.145
average fetch time (msec)	15.015	14.987	15.805	17.482	17.477	17.701
average disk utilization	0.82	0.38	0.25	0.21	0.17	0.14
Fetch time 15						
fetches	4698	5803	6051	4044	3922	3872
driver time (sec)	2.349	2.9015	3.0255	2.022	1.961	1.936
stall time (sec)	3.994	0.153	0.258	0	0	0.001
elapsed time (sec)	85.559	82.27	82.499	81.238	81.177	81.153
average fetch time (msec)	15.033	16.514	15.716	16.161	16.335	17.349
average disk utilization	0.83	0.58	0.38	0.2	0.16	0.14
Fetch time 30						
fetches	4698	5833	6194	6127	6200	5032
driver time (sec)	2.349	2.9165	3.097	3.0635	3.1	2.516
stall time (sec)	3.994	0.153	0.258	0	0	0.001
elapsed time (sec)	85.559	82.285	82.571	82.279	82.316	81.733
average fetch time (msec)	15.032	16.567	15.956	16.585	16.583	17.107
average disk utilization	0.83	0.59	0.4	0.31	0.25	0.18
Fetch time 60						
fetches	4698	5837	6224	6160	6042	5910
driver time (sec)	2.349	2.9185	3.112	3.08	3.021	2.955
stall time (sec)	3.994	0.153	0.258	0	0	0.001
elapsed time (sec)	85.559	82.287	82.586	82.296	82.237	82.172
average fetch time (msec)	15.03	16.577	15.937	16.598	16.713	17.145
average disk utilization	0.83	0.59	0.4	0.31	0.25	0.21

Table 66: Forestall performance as a function of static fetch time estimate on the postgres-join trace.

Disks	1	2	3	4	5	6	7	8	10	12	16
Fetch time 2											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.691	11.831	4.806	1.562	0.524	0.143	0.032	0.009	0.005	0.001	0
elapsed time (sec)	43.711	24.851	17.826	14.582	13.544	13.163	13.052	13.029	13.025	13.021	13.02
average fetch time (msec)	13.985	14.783	14.672	15.051	14.995	15.409	15.265	15.549	15.372	15.21	15.114
average disk utilization	0.99	0.92	0.85	0.8	0.68	0.6	0.52	0.46	0.36	0.3	0.22
Fetch time 4											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.691	10.837	5.617	1.562	0.524	0.143	0.032	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.857	18.637	14.582	13.544	13.163	13.052	13.029	13.025	13.021	13.02
average fetch time (msec)	13.985	14.13	14.653	15.051	14.995	15.409	15.265	15.549	15.372	15.21	15.114
average disk utilization	0.99	0.91	0.81	0.8	0.68	0.6	0.52	0.46	0.36	0.3	0.22
Fetch time 8											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.691	10.811	3.517	0.868	0.736	0.08	0.032	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.831	16.537	13.888	13.756	13.1	13.052	13.029	13.025	13.021	13.02
average fetch time (msec)	13.985	14.164	13.915	14.467	14.835	15.339	15.267	15.549	15.372	15.21	15.114
average disk utilization	0.99	0.92	0.87	0.8	0.67	0.6	0.52	0.46	0.36	0.3	0.22
Fetch time 15											
fetches	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425	1.5425
stall time (sec)	30.691	10.791	3.517	0.844	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.811	16.537	13.864	13.02	13.021	13.02	13.029	13.025	13.021	13.02
average fetch time (msec)	13.985	14.166	13.936	14.514	14.38	15.096	14.868	15.383	15.289	15.182	15.132
average disk utilization	0.99	0.92	0.87	0.81	0.68	0.6	0.5	0.46	0.36	0.3	0.22
Fetch time 30											
fetches	3085	3085	3085	3085	3085	3297	3764	3588	3123	3085	3085
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.5425	1.6485	1.882	1.794	1.5615	1.5425	1.5425
stall time (sec)	30.691	10.772	3.517	0.844	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.792	16.537	13.864	13.02	13.127	13.36	13.281	13.044	13.021	13.02
average fetch time (msec)	13.985	14.172	13.947	14.536	14.398	15.058	15.214	15.234	15.241	15.096	15.052
average disk utilization	0.99	0.92	0.87	0.81	0.68	0.63	0.61	0.51	0.36	0.3	0.22
Fetch time 60											
fetches	3085	3085	3085	3085	3166	3313	3830	3942	3904	3860	4108
driver time (sec)	1.5425	1.5425	1.5425	1.5425	1.583	1.6565	1.915	1.971	1.952	1.93	2.054
stall time (sec)	30.691	10.772	3.517	0.844	0	0.001	0	0.009	0.005	0.001	0
elapsed time (sec)	43.711	23.792	16.537	13.864	13.061	13.135	13.393	13.458	13.435	13.409	13.532
average fetch time (msec)	13.985	14.173	13.95	14.548	14.243	15.037	15.228	15.284	14.8	14.764	14.865
average disk utilization	0.99	0.92	0.87	0.81	0.69	0.63	0.62	0.56	0.43	0.35	0.28

Table 67: Fore stall performance as a function of static fetch time estimate on the postgres-select trace.

Disks	1	2	3	4	5	6
Fetch time 2						
fetches	5925	5912	5889	5891	5889	5897
driver time (sec)	2.9625	2.956	2.9445	2.9455	2.9445	2.9485
stall time (sec)	30.831	3.009	3.435	1.188	0.52	0.1
elapsed time (sec)	63.872	36.044	36.458	34.212	33.543	33.127
average fetch time (msec)	10.717	7.708	14.253	10.062	15.601	11.066
average disk utilization	0.99	0.63	0.77	0.43	0.55	0.33
Fetch time 4						
fetches	5925	6016	5907	5894	5890	5897
driver time (sec)	2.9625	3.008	2.9535	2.947	2.945	2.9485
stall time (sec)	30.667	0.421	3.216	0.707	0.546	0.1
elapsed time (sec)	63.708	33.508	36.248	33.733	33.57	33.127
average fetch time (msec)	10.711	7.706	14.132	9.903	15.673	11.06
average disk utilization	1	0.69	0.77	0.43	0.55	0.33
Fetch time 8						
fetches	5925	7045	6444	6025	5910	5896
driver time (sec)	2.9625	3.5225	3.222	3.0125	2.955	2.948
stall time (sec)	30.667	0.337	0.355	0.16	0.225	0.06
elapsed time (sec)	63.708	33.938	33.656	33.251	33.259	33.087
average fetch time (msec)	10.711	7.499	14.085	10.019	14.989	10.639
average disk utilization	1	0.78	0.9	0.45	0.53	0.32
Fetch time 15						
fetches	5925	7274	6563	8643	7534	6421
driver time (sec)	2.9625	3.637	3.2815	4.3215	3.767	3.2105
stall time (sec)	30.667	0.337	0.356	0.129	0.133	0.055
elapsed time (sec)	63.708	34.053	33.716	34.529	33.979	33.344
average fetch time (msec)	10.711	7.492	14.095	9.717	15.34	10.458
average disk utilization	1	0.8	0.91	0.61	0.68	0.34
Fetch time 30						
fetches	5925	7742	6563	9699	8300	9846
driver time (sec)	2.9625	3.871	3.2815	4.8495	4.15	4.923
stall time (sec)	30.667	0.337	0.356	0.129	0.133	0.055
elapsed time (sec)	63.708	34.287	33.716	35.057	34.362	35.057
average fetch time (msec)	10.711	7.467	14.099	9.758	15.445	10.708
average disk utilization	1	0.84	0.91	0.67	0.75	0.5
Fetch time 60						
fetches	5925	7778	6563	9807	8300	10015
driver time (sec)	2.9625	3.889	3.2815	4.9035	4.15	5.0075
stall time (sec)	30.667	0.337	0.356	0.129	0.133	0.055
elapsed time (sec)	63.708	34.305	33.716	35.111	34.362	35.141
average fetch time (msec)	10.711	7.497	14.096	9.798	15.451	10.678
average disk utilization	1	0.85	0.91	0.68	0.75	0.51

Table 68: Forestall performance as a function of static fetch time estimate on the xds trace.