An Analysis of Communication Induced Checkpointing

Lorenzo Alvisi *‡ Elmootazbellah Elnozahy† Sriram Rao* ‡ Syed Amir Husain‡ Asanka De Mel‡

‡ Department of Computer Sciences
U. T. Austin

†IBM Austin Research Lab
Austin, Texas, USA.

Abstract

Communication induced checkpointing (CIC) allows processes in a distributed computation to take independent checkpoints and to avoid the domino effect. This paper presents an analysis of CIC protocols based on a prototype implementation and validated simulations. Our result inidcate that there is sufficient evidence to suspect that much of the conventional wisdom about these protocols is questionable.

1 Introduction

There are three styles for implementing application-transparent rollback-recovery in message-passing systems, namely coordinated checkpointing, message logging, and communication-induced checkpointing (CIC) [5]. Both coordinated checkpointing and message logging have received considerable analysis in the literature [6, 12, 14, 15], but little is known about the behavior of CIC protocols. This paper presents an experimental analysis of these protocols through a prototype implementation and reveals several of their theoretical and pragmatic characteristics.

CIC protocols are believed to have several advantages over other styles of rollback-recovery. For instance, they allow processes considerable autonomy in deciding when to take checkpoints. A process can thus take a checkpoint at times when saving the state would incur a small overhead [10, 13]. CIC protocols also are believed to scale up well with a larger number of processes since they do not require the processes to participate in a global checkpoint. But there is a price to pay for these advantages. First, the protocol-specific information piggybacked on application messages occasionally "induces" processes to take forced checkpoints before they can process the messages. Second, processes have to pay the overhead of piggybacking information on top of application messages, and they also need to keep several checkpoints on stable storage. These advantages and disadvantages are clearly qualitative and potentially arguable. A purpose of our work is to shed some light on these issues using quantitative metrics drawn from a real system.

To this end, we have implemented three CIC protocols—-the original one by Briatico et al [2], and two recent protocols [1, 7] based on the Z-path theory [11]— and we have examined their performance using two metrics, namely the average number of forced checkpoints a protocol causes and its running time. The first metric is important because forced checkpoints negate the autonomy advantage of CIC protocols. Also, they contribute substantially to the performance and resource overheads. A good CIC protocol therefore tries to limit these forced checkpoints to the extent possible. The experiments use four compute-intensive programs from the NPB 2.3 benchmark suite [3], which is a representative of a class of applications that have traditionally been the primary users of checkpointing protocols. We then use the implementation in part to validate a simulation model that we built to study further the scalability of the protocols and their behaviors under different communication patterns.

Our results reveal several important properties of CIC protocols and highlight several implementation and theoretical issues that were not addressed in the literature. We hope that our work will be a first step in investigating an area that thus far has been only subject to theoretical research.

*This author supported in part by the National Science Foundation (CAREER award CCR-9734185, Research Infrastructure Award CDA-9734185, and DARPA/SPAWAR grant N66001-98-8911).

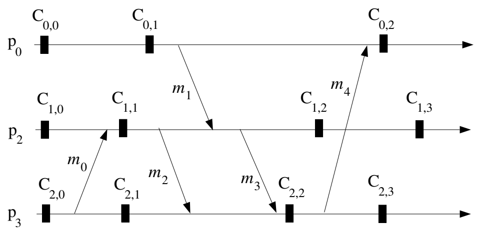

Figure 1 : A distributed computation. $C_{i,j}$ denotes the j^{th} checkpoint of process p_i.

2 Background

This section reviews the three protocols used in the experiments, along with some necessary definitions. The description is inevitably terse and covers only the features necessary to follow the experimental work described later.

2.1 Definitions

Local checkpoints: A process may take a local checkpoint any time during the execution. The local checkpoints of different processes are not coordinated to form a global consistent checkpoint [4].
Forced checkpoints: To guard against the domino effect, a CIC protocol piggybacks protocol-specific information to application messages that processes exchange. Each process examines the information and occasionally is forced to take a checkpoint according to the protocol.
Useless checkpoints: A useless checkpoint of a process is one that will never be part of a global consistent state [18]. In Figure 1, checkpoint $C_{2,2}$ is an example of a useless checkpoint. Useless checkpoints are not desirable because they do not contribute to the recovery of the system from failures, but they consume resources and cause performance overhead.

Checkpoint intervals: A checkpoint interval is the sequence of events between two consecutive checkpoints in the execution of a process.

2.2 Z-paths and Z-cycles

Z-paths: A Z-path (zigzag path) is a special sequence of messages that connects two checkpoints. Let $\rightarrow$ denote Lamport's happen-before relation [9]. Given two local checkpoints $C_{i,m}$ and $C_{j,n}$, a Z-path exists between $C_{i,m}$ and $C_{j,n}$ if and only if one of the following two conditions holds:

1. $m < n$ and $i = j$; or
2. There exists a sequence of messages $[m_0, m_1, , m_z]$, $z \geq 0$, such that:
   (a) $C_{i,m} \rightarrow$ $send_i(m_0)$;
   (b) $\forall l < z$, either $deliver_k(m_l)$ and $send_k(m_{l+1})$ are in the same checkpoint interval, or $deliver_k(m_l) \rightarrow send_k(m_{l+1})$; and
   (c) $deliver_j(m_z) \rightarrow C_{j,n}$ where $send_i$ and $deliver_i$ are communication events executed by process p_i. In Figure 1, $[m_1, m_2]$ and $[m_1, m_3]$ are examples of Z-paths.

Z-cycles: A Z-cycle is a Z-path that begins and ends with the same checkpoint. In Figure 1, the Z-path $[m_4, m_1, m_3]$ is a Z-cycle that involves checkpoint $C_{2,2}$.

2.3 Z-cycles and CIC

CIC protocols do not take useless checkpoints. These protocols recognize that the creation of useless checkpoints depends on the occurrence of specific patterns in which processes communicate and take checkpoints [8]. Informally, these protocols recognize potentially dangerous patterns and break them before they occur. This intuition has been formalized in an elegant theory based on the notion Z-cycles. A key result in this theory is that a local checkpoint is useless if it is involved in a Z-cycle [8, 11]. Hence, to avoid useless checkpoints it suffices that no Z-path ever becomes a Z-cycle. Enforcing the no-Z-cycle ($\mathcal{NZC}$) condition may require that a process save additional forced checkpoints in addition to its local checkpoints. There are two approaches to avoiding Z-cycles. The first approach uses a function that associates a timestamp with each checkpoint. The protocol guarantees, through forced checkpoints if necessary, that (*i*) if there are two checkpoints $C_{i,m}$ and $C_{j,n}$ such that $C_{i,m} \rightarrow C_{j,n}$, then $ts(C_{j,n}) \geq ts(C_{i,m})$, where $ts(C)$ is the timestamp associated with checkpoint C; and (*ii*) consecutive local checkpoints of a process have increasing timestamps. The second approach relies instead on preventing the formation of specific checkpoint and communication patterns that may lead to the creation of a Z-cycle. Protocols that follow this approach do not adopt a specific function for associating timestamps with checkpoints. However, for these protocols there always exists an equivalent time-stamping function that would cause the same forced checkpoints [8].

2.4 Briatico, Ciuffoletti and Simoncini (BCS)

In BCS, each process p_i maintains a logical clock lc_i that functions as p_i's checkpoint timestamp. The timestamp is an integer variable with initial value 0 and is incremented according to the following function:

1. lc_i increases by 1 whenever p_i takes a local checkpoint.
2. p_i piggybacks on every message m it sends a copy of the current value of lc_i. We denote the piggybacked value as $m.lc$.
3. Whenever p_i receives a message m, it compares lc_i with $m.lc$. If $m.lc > lc_i$, then p_i sets lc_i to the value of $m.lc$ and takes a forced checkpoint before processing the message.

The set of checkpoints having the same timestamps in different processes is guaranteed to be a consistent state. Therefore, this protocol guarantees that there is always a recovery line corresponding to the lowest timestamp in the system, and the domino effect is prevented.

2.5 Hélary, Mostefaoui, Netzer and Raynal (HMNR)

The HMNR protocol uses the observation that if checkpoints' timestamps always increase along a Z-path (as opposed as simply non-decreasing, as required by rule (*i*) above), then no Z-cycle can ever form. It is thus possible to design functions that take advantage of this observation. Hélary et al start with the following simple scheme which would require each process to maintain a logical clock, as in BCS, and to apply the following rules:

1. lc_i increases by 1 whenever p_i takes a local or forced checkpoint.
2. Whenever p_i sends a message m, it piggybacks on m a copy of lc_i, and we denote this value by $m.lc$ as before.
3. Whenever p_i receives a message m, it compares lc_i with $m.lc$. If $m.lc > lc_i$, then p_i sets lc_i to the value of $m.lc$.

Then, they refine the protocol using more sophisticated observations and requiring that processes append more information. Figure 2, adapted from [7], shows how this simple timestamp function can be used to decide when to take a checkpoint. Let $ts(m)$ denote the timestamp piggybacked on message m. When process p_1 receives message m_0, if $ts(m_0) \leq ts(m_1)$ then certainly $ts(C_{0,0}) \leq ts(C_{2,1})$ and there is no need for a forced checkpoint. If $ts(m_0) > ts(m_1)$, however, delivering m_1 may create the possibility of generating a Z-path along which the timestamps do not increase. A straightforward way to avoid this risk is to force p_1 to take a checkpoint before delivering m_0, thereby breaking the Z-path. It may be possible, however, for p_1 to avoid taking a forced checkpoint by using a more sophisticated timestamp function. For instance, a function that piggybacks on application messages information about the logical clocks of all processes may give process p_1 more information to decide if a forced checkpoint is really necessary. In Figure 2 (b), if p_1 knows that the value of p_2's local clock is at least x when it is about to deliver m_0, then even if $ts(m_0) > ts(m_1)$, p_1 does not need to take a forced checkpoint if $ts(C_{0,0}) \leq ts(m_1) \leq x < ts(C_{2,1})$. HNMR uses this observation and more sophisticated ones to reduce the number of forced checkpoints while still ensuring that the timestamps always increase along a Z-path. In [7], Hélary et al present several CIC algorithms. For our experiments, we have used the most sophisticated one, which reduces as much as possible the number of forced checkpoints.

2.6 Baldoni, Quaglia and Ciciani (BQC)

The BQC protocol does not prevent forced checkpoints by using an explicit function to timestamp checkpoints [1]. Rather, this protocol enforces $\mathcal{NZC}$ by preventing the formation of patterns of checkpoints and communication that may result in the creation of

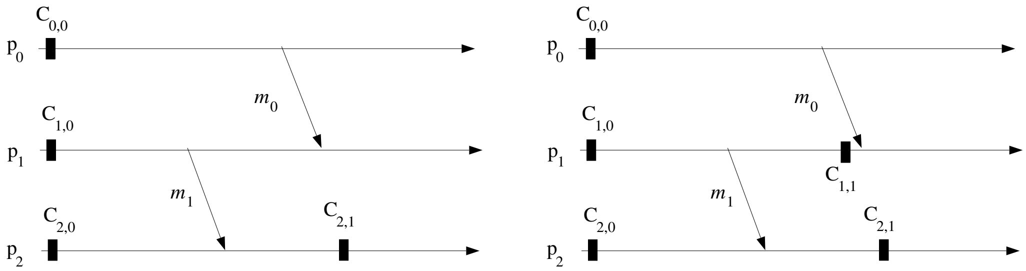

Figure 2 : To checkpoint or not to checkpoint?

a Z-cycle. In particular, BQC prevents the creation of suspected Z-cycles. Figure 3 shows the structure of a suspected Z-cycle. In this example, process p_i would take a checkpoint before delivering message m_2, in order to avoid a potential Z-cycle that includes m_0 and m_1 and involves checkpoint $C_{0,1}$. Note that a suspected Z-cycle is not necessarily a Z-cycle. For instance, there is no Z-cycle in Figure 3 unless there is an actual Z-path starting with m_0 and ending with m_1 between $C_{i,0}$ and $C_{0,1}$. This information may or may not be available to process p_i when it receives m_2. If the information is not available, the protocol opts for safety and takes a forced checkpoint before delivering m_2. If the information is available, however, the protocol refrains from taking a forced checkpoint. The actual protocol propagates n^2 values on each application message to help processes detect suspected Z-cycles and suppress them using forced checkpoints [1]. The reference also contains a formal characterization of the notion of suspected Z-cycle and a proof that a protocol that prevents suspected Z-cycle also satisfies $\mathcal{NZC}$ [1].

3 Implementation Issues

An implementation of CIC must deal with several pragmatic issues that are typically left out of protocol specifications. We describe how we resolved these issues in our implementation.

3.1 Autonomy in Local Checkpoints

A stated advantage of CIC protocols is that they prevent the domino effect while allowing processes considerable autonomy in deciding when to take local checkpoints. An efficient implementation of CIC must therefore adopt a checkpointing policy that exploits this autonomy and translates it into a benefit. In general, this requires a good understanding of the application and of the execution environment. For example, detailed knowledge of the application may allow processes to take checkpoints when the size of the live variables is small [10, 13].

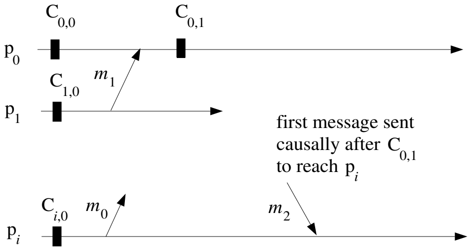

Figure 3 : A suspect Z-cycle involving checkpoint $C_{0,1}$.

In our study, we have chosen a set of compute-intensive, long-running applications that have traditionally been the main beneficiary of checkpointing protocols. We have found that either the complexity of the application program precludes investing the effort in defining a reasonable checkpoint placement policy, or that the application structure does not reveal points within the execution where taking a checkpoint is more advantageous than others. Therefore, we have resolved to use a probabilistic distribution to emulate what an autonomous application would do in deciding on where to place the local checkpoints. In addition, we have also used the traditional policy of taking checkpoints at regular intervals. The results were similar; indeed, one of the first conclusions that we reached in our implementation is that even with a deep understanding of the applications' structure, no policy for taking local checkpoints can reasonably be adopted without considering the effect of the forced checkpoints that occur because of communication events. For example, a pre-run analysis may decide on the execution points during which it is most "convenient" to schedule local checkpoints within a particular process. But if this policy ignores the forced checkpoints, then it is possible to schedule a local checkpoint immediately after a forced checkpoint. In this case, taking the local checkpoint will result in additional work and overhead, with no substantial reduction of the amount of work at risk. A more reasonable decision may then be to postpone taking the local checkpoint. In any case, the autonomy of deciding when to take checkpoints seems to be limited by the occurrence of forced checkpoints due to interactions with other processes.

3.2 Non-blocking Checkpointing in CIC

The benefits of non-blocking checkpointing in reducing the performance overhead of checkpointing protocols have been clearly established [6, 14]. Non-blocking checkpointing allows the application to resume computation as soon as possible and to schedule the actual writing of the checkpoint concurrently with the application execution. The result is that saving the checkpoint to stable storage does not become a bottleneck that impedes application progress. However, using non-blocking checkpointing in CIC introduces potential inconsistencies with the specification of the protocols. To understand why, consider the situation where a process p_i takes a non-blocking checkpoint (local or forced) and then sends a message m to another process p_j; assume furthermore that the CIC protocol being used requires p_j to take a forced checkpoint before delivering m. Recall that all CIC protocols maintain the invariant that no Z-cycles may form and no useless checkpoint is ever taken: hence, the forced checkpoint of p_j should never be useless and no Z-cycles can include it. However, suppose that p_i fails after sending m, but before p_i's checkpoint has been entirely written to stable storage. The failure of p_i makes the checkpoint taken by p_j useless, thereby violating the protocol invariant. The problem is not just cosmetic, because several such communication events may occur while several non-blocking checkpoints are being written to stable storage according to an implementation of non-blocking checkpointing. Therefore, it may be possible for many checkpoints on stable storage to be rendered useless because a failure of some process in the system occurred before one or more of its checkpoints were saved on stable storage. Indeed, in situations like this, Z-cycles do form and useless checkpoints are taken. There are two solutions to this problem. The first one blocks any outgoing messages from a process until all its non-blocking checkpoints have been written to stable storage. The messages are buffered and released only when a process receives notification from the checkpointing agent that the checkpoint has been saved. There is a penalty to pay for this modification, but it is preferable to disallowing non-blocking checkpointing altogether. It is interesting to note here that this solution shows how a pragmatic consideration may require a modification in the protocol itself to maintain its invariant. The second solution that we considered is to simply allow these temporary Z-cycles to form, and hope that the checkpoints will be written before a failure actually occurs. In a sense, this is an optimistic implementation of CIC, which may allow Z-cycles to form temporarily while some checkpoints are being written to stable storage. If the optimistic assumption holds, then the invariant of the protocol is preserved and no useless checkpoints are ever taken. However, if the assumption is violated because of a failure, then some of the checkpoints that would have otherwise be part of the recovery line will have to be discarded. The benefit of this optimistic alternative is that the overhead is small and does not require any modification to the protocol as specified. After an interesting debate among the authors, it was resolved to use the second solution for its simplicity and because failures are supposedly rare. It is important however for future implementers to understand the subtle issues involved with the pragmatic choices described here and how they may affect the protocol implementation.

4 Experiments and Analysis

The testbed for this study consists of four 300-MHz Pentium-II based workstations connected by a 100MB/s Ethernet. Each workstation has two processors, 512MB of RAM, and a 4GB disk used to implement stable storage. The machines ran Solaris 2.6, and used Suns f77 and C compilers. The testbed is part of the Egida tool [16], which includes support for incremental checkpointing and implements non-blocking checkpointing by forking off a child process that writes the checkpoint to stable storage. The applications under study consist of four MPI [17] programs from the NPB 2.3 benchmark suite [3]. These programs represent common computational loads in fluid dynamics applications and typify the kind of applications that have traditionally benefited from checkpointing; their characteristics are given in Table 1.The performance metrics we report are the number of forced checkpoints that a protocol causes and the performance overhead. We use a combination of experiments on the prototype implementation, and then we use the prototype itself to validate a simulator that we built to study CIC protocol further, under different communication patterns and environments.

4.1 The Measured Performance of CIC Protocols

The first set of experiments consists of running each of the four applications under the three protocols and for two local checkpoint placement policies. The first policy triggers local checkpoints according to an exponential distribution with a mean checkpoint interval (μ) set to 360 seconds, while the second policy uses the same probabilistic distribution but with the mean checkpoint interval set to 480 seconds. Table 2 shows the results of the experiments. It reports the execution time of the entire application, in addition to the per-process average number of local and forced checkpoints. For convenience, the average per-process total number of checkpoints is also reported. The table also reports the per-process average checkpoint size (either local or forced).

Analysis The results reveal a few issues. In BCS and HMNR, the number of forced checkpoints is essentially the same. In contrast, the BQC protocol is showing a comparatively larger number of forced checkpoints when compared to the other two.. The reason for this can be ascribed to the communication pattern of the applications under study. These applications use a common iterative structure to solve a computationally intensive problem in which processes exchange partial results and resume. This leads to a communication pattern that mimics a periodic broadcast. Under this pattern, the BQC protocol seems to be too "eager" in preventing Z-cycles compared to BCS and HMNR. Furthermore, two additional effects seem to occur:

1. Many suspected Z-cycles end up causing forced checkpoints without actually being a menace.
2. It is often the case that more than one process "volunteer" in parallel to break the same suspected Z-cycle by forcing checkpoints.

Consider Figure 4. In this example, we see process p_2 take a forced checkpoint because of message m_3, not knowing that process p_1 has already broken the suspected Z-cycle (part (2) of the figure). Similarly, process p_0 takes a forced checkpoint because of message m_2, not knowing that process p_1 has already broken the Z-cycle using checkpoint $C_{1,1}$. This behavior continues for a while under the communication pattern used by our applications. This suggests that there is a disadvantage to using CIC protocols that suspect Z-cycles or that are eager to prevent a Z-cycle from forming before it is actually clear that one is indeed forming. In contrast, protocols that use time-stamping functions adopt a lazy approach: they prevent the formation of Z-cycles only at the last possible moment, and therefore work better. Additionally, the results show that the alleged benefit of process autonomy in placing local checkpoints does not materialize in practice. Under the best circumstances, a process takes twice as many forced checkpoints as local ones. The curious notion of process autonomy in distributed systems where all processes become inter-dependent seems to be on shaky ground. The results also point out to another serious problem with CIC protocols in general, namely the unpredictability of the checkpointing rate. In all experiments, the protocols ended up taking more checkpoints than could be anticipated based on the local distribution of checkpoint placement. For BCS and HMNR, the number of forced checkpoints was generally twice the number of local ones. For BQC, the ratio was worse. The ratio in itself is a function of the application, the number of processes, and the checkpoint placement. The fact that it is unpredictable makes the protocols cumbersome to use in practice, because it is difficult to plan ahead of time the actual stable storage requirements and the mean checkpointing interval. Contrast this with consistent checkpointing protocols where the number of checkpoints and required stable storage can be estimated with great certainty beforehand [6]. The table also points to another negative aspect of using CIC protocols. The performance overhead when considering the running time was relatively bad, reaching between 5 % to

Application	NPB Specific Info.	Messages/sec.	Message Size (Avg.) (KB)	Exec. Time (sec.)
bt	Class A	6	50.5	1530
cg	Class B	14	60.7	1516
lu	Class A	54	3.7	975
sp	Class A	17	43.4	1222

Table 1: Characteristics of the benchmarks used in the experiments.

Application	Protocols	μ	Number of Checkpoints Local	Forced	Total	Ckpt Size (MB)	Exec.Time (sec.)
bt	BCS	360	6	13	19	69.5	1777
		480	4	9	13	70.9	1715
	HMNR	360	5	11	16	70.5	1709
		480	4	9	13	71.0	1683
	BQC	360	6	54	60	41.5	1875
		480	4	39	43	41.0	1819
cg	BCS	360	5	11	16	17.6	1683
		480	4	9	13	18.9	1655
	HMNR	360	5	11	16	17.7	1655
		480	3	10	13	19.1	1643
	BQC	360	5	24	29	8.5	1689
		480	3	16	19	13.5	1665
lu	BCS	360	4	9	13	10.7	1051
		480	2	4	6	11.0	1033
	HMNR	360	4	9	13	10.8	1035
		480	2	4	6	11.0	1015
	BQC	360	4	33	37	6.6	1050
		480	2	14	16	6.0	1036
sp	BCS	360	5	11	16	20.8	1339
		480	3	7	10	21.2	1290
	HMNR	360	5	11	16	20.7	1320
		480	3	7	10	21.1	1300
	BQC	360	5	37	42	11.8	1362
		480	3	31	34	12.2	1329

Table 2: Performance of three CIC protocols for two checkpoint intervals and four applications.

20 % of the execution time. This behavior is actually common in systems where checkpoints are not coordinated and processes communicate frequently [6]. In these situations, when a process takes a local checkpoint independent of the others it inevitably slows down because of the state saving and memory copying that occur during the checkpoint. This in turn delays the production of the expected partial result that the process will send to others in the next communication round. Consequently, the slowdown affects other processes even if they are not taking checkpoints in the meantime. The resulting slowdowns stagger quickly and have a cumulative effect because many of these independent checkpoints occur at different times [6]. Finally, we would like to note that incremental checkpointing seems to mitigate some of the effects of having to take so many checkpoints (forced or local). The results show that the average per-process checkpoint size goes down as the frequency of checkpointing increases, just as one would expect. In summary, lazy protocols that use time-stamping to break Z-cycles perform better than eager protocols that take forced checkpoints as soon as they suspect a Z-cycle. The unpredictability of the actual number of checkpoints to be taken (forced and local) makes these protocols cumbersome to use in practice because no reasonable planning of resources and checkpointing frequency can be made without understanding the application and its communication patterns. Also, it seems that any hope of a benefit for allowing processes to take independent checkpoints is thwarted by the fact that a process ends up taking at least twice as many forced checkpoints than local ones. And finally, CIC protocols share some of the negative performance properties of independent checkpointing when used in computations where the processes are tightly coupled and communicate frequently.

4.2 Scalability and Effects of Communication Patterns

To assess the effect of increasing the number of processes on the protocol performance, we constructed a simulator to measure the number of forced checkpoints for each of the three protocols. We first validated the simulator using the measured number of forced checkpoints for 4 processes, and we then used it to estimate the number of forced checkpoints for different numbers of processors and different communication patterns. Figure 5 shows the results for the three protocols as the number of processes in the computation varies. Two different sets of measurements are reported. The first set is for a random distribution of messages with a relatively low load, where each process sends an average of 10 messages between each two consecutive local checkpoints. That is, in this simulation processes do not communicate much and communicate

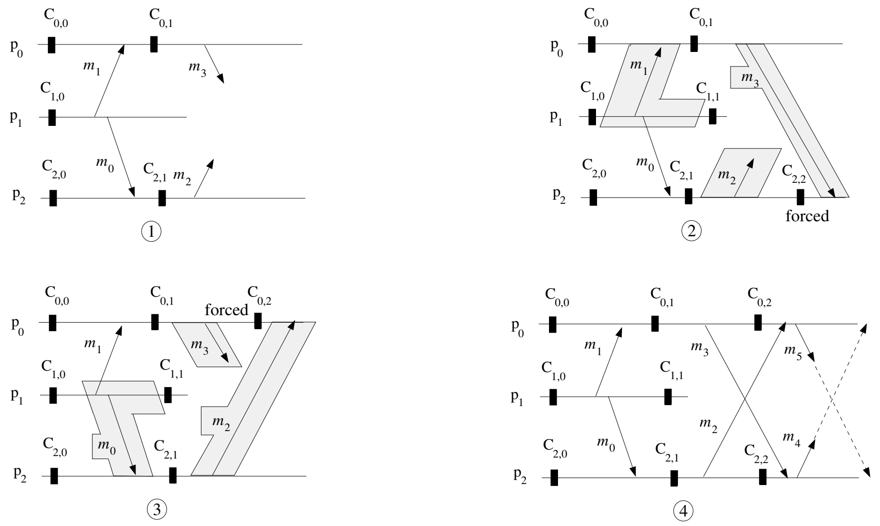

Figure 4 : Anomalies in detecting suspect Z-cycles.

with different processes equally at random and at random intervals. During this simulation, 119 local checkpoints were taken on average. The second set of measurements shows the same results but with a different communication pattern in which each process talks to two designated neighbors at uniform intervals. A process sends about 500 messages between each two consecutive local checkpoints. This communication pattern is representative of those that occur in distributed over-relaxation algorithms.

Analysis The results show that, in general, the communication pattern strongly affects the behavior of CIC protocols. This is expected. But the results also show that CIC protocols do not scale well. In both cases, there is an almost linear increase in the number of forced checkpoints per process as the number of processes increase. For this set of applications, at least, the conventional wisdom that these protocols scale better because they do not resort to global coordination is not true. The results also show that CIC protocols seem to favor random patterns of communications with low loads.

4.3 An Adaptive Local Checkpointing Policy

The results of the experiments so far suggest that a flurry of forced checkpoints occur throughout the system as a result of one process taking a local checkpoint. It is plausible that if forced checkpoints are not taken into account, a local checkpointing policy may take a local checkpoint shortly after a forced checkpoint has been taken. Such a local checkpoint advances the recovery point of the process by a very short amount compared to the previous forced checkpoint. Furthermore, this local checkpoint is likely to trigger more forced checkpoints in other processes, escalating the phenomenon even further. It may be argued that the resulting overhead can be limited by using incremental checkpointing, and that therefore the local checkpoint does not have to save a lot of state on stable storage if a forced checkpoint has been taken recently. But we contend that taking a checkpoint, however small, always has an overhead associated with it, if only to compute the state that must be saved and to arrange for the copy-on-write to implement non-blocking checkpointing. However it may be, this overhead cannot be ignored. Therefore, there is very little to gain by taking this local checkpoint, while there is a potential for larger overhead. To fix this problem, we experimented with an adaptive local checkpointing policy that refrains from taking a scheduled local checkpoint if a forced checkpoint has occurred during the last T seconds, where T is a parameter. Figure 6 shows the resulting number of local and forced checkpoints for the four applications and the three protocols under study. We report three measurements, one with the adaptive policy disabled (T = 0) and two for different values of T (60 and 90 seconds). For each T, the table shows the number of local and forced checkpoints under each of the three protocols. The measurements for different applications are reported separately.

Analysis The results show that taking forced checkpointing into account reduces the number of local checkpoints, which in turn reduces the number of forced checkpoints. The results are more pronounced for the BQC protocol. These results show two things:

1. A successful local checkpoint placement policy must contain a dynamic element that accounts for the occurrence of forced checkpoints. The simple policy "let us checkpoint every x seconds" does not work well.
2. A successful local checkpoint placement policy must adapt to the application communication patterns as they change during execution. This would reduce the frequency of checkpoints when the communication load is heavy and the frequency of forced checkpoints is high, and vice versa.

Our recommendations once more outline the unpredictability that faces a user of CIC protocols in practice, though they outline plausible solutions. It is perhaps possible to come up with better placement policies than the one we outlined here, but this is out of the paper's scope.

5 Conclusions

We have conducted several experiments to analyze the behavior and characteristics of communication-induced checkpointing for a class of compute intensive distributed applications. Our results show that:

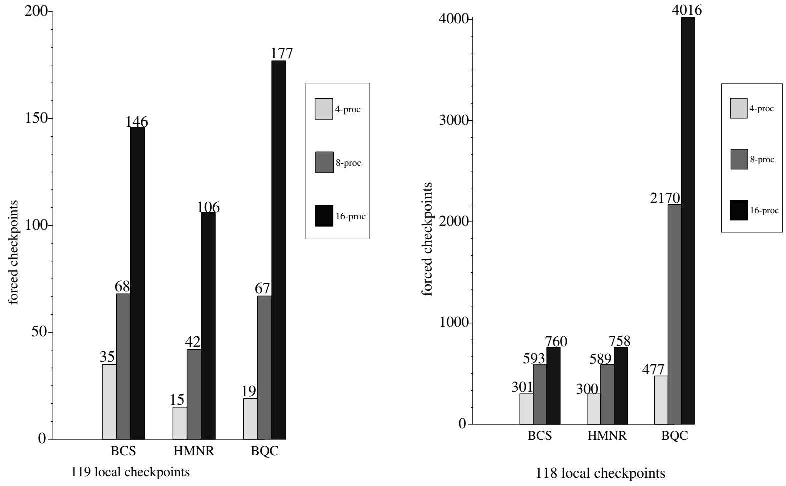

(a) (b)

Figure 5 : The effect of adaptive local checkpointing: Number of local and forced checkpoints for the three protocols under values of T for the four applications under study.

1. CIC protocols that use an eager approach to preventing Z-cycles by taking forced checkpoints whenever they suspect the formation of a Z-cycle are bound to perform worse than lazy protocols that use a time-stamping function to prevent a Z-cycle at the last possible second.

2. CIC protocols do not scale well with a larger number of processes. We have found that the number of forced checkpoints increases almost linearly with the number of processes.

3. A process takes at least twice as many forced checkpoints as local ones. Therefore, the touted benefit of autonomy of CIC protocols in allowing the processes to take independent checkpoints does not seem to materialize in practice.

4. There is a considerable unpredictability in the way CIC protocols behave in practice. The amount of stable storage required, performance overhead, and number of forced checkpoints depend greatly on the number of processes, the application, and the communication pattern. This unpredictability makes the use of CIC protocols in practice more cumbersome than other alternatives.

5. A successful placement policy of local checkpoints must be dynamic, must account for forced checkpoints, and must adapt to changes in the application behavior.

6. CIC protocols seem to perform best when the communication load is low and the pattern is random. Regular, heavy load communication patterns seem to fare worse.

We would like to stress that the results are only valid for the application set that we have studied—we lay no claim that these results generalize to all applications. Nevertheless, we believe that there is sufficient evidence to suspect that much of the conventional wisdom about these protocols is questionable and that more experimental work to investigate these protocols further.

References

[1] R. Baldoni, F. Quaglia, and B. Ciciani. A VP-Accordant Checkpointing Protocol Preventing Useless Checkpoints. In *Proceedings of International Symposium on Reliable Distributed Systems*, pages 61–67, West Lafayette, IN, October 1998.

[2] D. Briatico, A. Ciuffoletti, and L. Simoncini. A Distributed Domino-Effect Free Recovery Algorithm. In *Proceedings of the IEEE International Symposium on Reliability Distributed Software and Database*, pages 207–215, December 1984.

[3] NASA Ames Research Center. NAS Parallel Benchmarks. http://science.nas.nasa.gov/Software/NPB/, 1997.

[4] K. M. Chandy and L. Lamport. Distributed snapshots: determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, February 1985.

[5] E. N. Elnozahy, D. B. Johnson, and Y. M. Wang. A Survey of Rollback-Recovery Protocols in Message-Passing Systems. Technical Report CMU-CS-96-181, Carnegie Mellon University, 1996.

[6] E. N. Elnozahy, D. B. Johnson, and W. Zwaenepoel. The Performance of Consistent Checkpointing. In *Proceedings of the Eleventh Symposium on Reliable Distributed Systems*, pages 39–47, October 1992.

[7] J. M. Hélary, A. Mostefaoui, R. H. B. Netzer, and M. Raynal. Preventing Useless Checkpoints in Distributed Computations. In *Proceedings of IEEE International Symposium on Reliable Distributed Systems*, pages 183–190, 1997.

[8] J. M. Hélary, A. Mostefaoui, and M. Raynal. Virtual precedence in asynchronous systems: concepts and applications. In *Proceedings of the 11th Workshop on Distributed Algorithms*. LNCS press, 1997.

[9] L. Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *Communications of the ACM*, 21(7):558–565, July 1978.

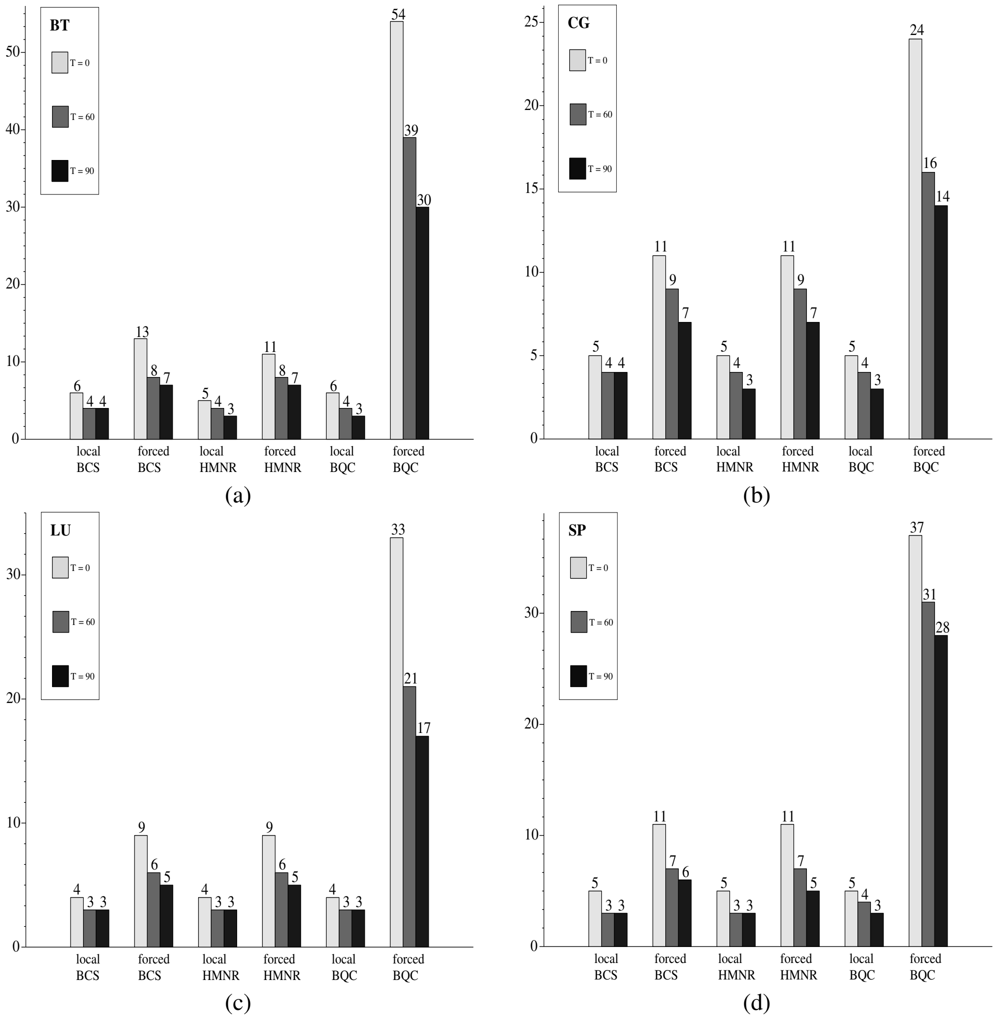

Figure 6 : The effect of adaptive local checkpointing: Number of local and forced checkpoints for the three protocols under values of T for the four applications under study.

[10] C. C. Li and W. K. Fuchs. CATCH: Compiler-assisted techniques for checkpointing. In *Proceedings of the 20th International Symposium on Fault-Tolerant Computing*, pages 74–81, 1990.

[11] R. Netzer and J. Xu. Necessary and Sufficient Conditions for Consistent Global Snapshots. Technical Report 93-32, Department of Computer Sciences, Brown University, July 1993.

[12] N. Neves and W. K. Fuchs. RENEW: A Tool for Fast and Efficient Implementation of Checkpoint Protocols. In *Proceedings of the 28th IEEE Fault-Tolerant Computing Symposium (FTCS)*, Munich, Germany, June 1998.

[13] J. S. Plank, M. Beck, and G. Kingsley. Compiler-assisted memory exclusion for fast checkpointing. *IEEE Technical Committee on Operating Systems Newsletter, Special Issue on Fault Tolerance*, pages 62—67, December 1995.

[14] J. S. Plank, M. Beck, G. Kingsley, and K. Li. Libckpt:Transparent checkpointing under Unix. In *Proceedings of the USENIX Technical Conference*, pages 213–224, January 1995.

[15] J. S. Plank and K. Li. Faster checkpointing with $N+1$ parity. In *Proceedings of the Twenty Fourth International Symposium on Fault-tolerant Computing (FTCS-24)*, pages 288—297, June 1994.

[16] S. Rao, L. Alvisi, and H. M. Vin. Egida: An Extensible Toolkit for Low-overhead Fault-tolerance. In *Proceedings of the IEEE Fault-Tolerant Computing Symposium (FTCS-29)*, Madison, WI, June 1999.

[17] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI: The Complete Reference*. Scientific and Engineering Computation Series. The MIT Press, Cambridge, MA, 1996.

[18] Y. M. Wang and W. K. Fuchs. Optimisitic message logging for independent checkpointing in message-passing systems. In *Proceedings of the 11th Symposium on Reliable Distributed Systems*, pages 147–154, October 1992.