Discontinuous Incremental: A New Approach Towards Extremely Lightweight Checkpoints

Viet Hai Ha, Eric Renault

Discontinuous Incremental: A new approach towards extremely lightweight checkpoints

Viet Hai Ha[1] and Éric Renault[2]
Institut Télécom - Télécom SudParis
Samovar UMR INT-CNRS 5157
Évry, France

Abstract—**Checkpointing is an important method for providing fault tolerance, load balancing, process migration, periodic backup, and many other functions [9], [14]. It is also the basic tool used in CAPE [1], [2], a paradigm which aims at distributing the execution of a program on a distributed-memory environment. This paper presents the new approach to checkpointer and the original optimization on checkpoint structure we have implemented and evaluated to make incremental checkpointing more efficient and more appropriate, especially for CAPE.**

I. INTRODUCTION

As a critical component to ensure a program will terminate its execution even in case of system or hardware failure, many works have been done to develop checkpointing techniques and different approaches have been used. The first solution named complete checkpointing consists in saving all the information related to the running process [4]–[7], regardless these information have been modified from the beginning of the execution or they can be retrieved easily for example from a dynamic shared library. The alternative to complete checkpoints are incremental checkpoints. In this case, only the information that have been modified since the beginning of the execution are effectively saved [8], [9]. The identification of the memory areas that have to be saved may either be provided by the developer of the application using pragma directives and/or dedicated functions, or automatically detected by the checkpointing tool itself. Both complete and incremental solutions have advantages and drawbacks. The main advantage of a complete checkpoint is the simplicity to generate it, while the main advantage of an incremental checkpoint is its size. An important drawback of incremental checkpointing technique, as a consequence of the regular monitoring of the process memory, is that the execution speed may significantly decrease. Overcoming this disadvantage is an important requirement to increase the performance of an incremental checkpointer.

OpenMP (for Open Multi-Processing) [15] is a very simple and powerful set of directives and functions to generate parallel programs from C, C++ or Fortran codes. The main limitation of OpenMP is that it is limited to shared-memory architectures. Some attempts have tried to port OpenMP on distributed memory architectures with various success. In order to achieve this goal, we have developed a new parallel computing paradigm called CAPE which aims at using checkpoints to distribute the execution of a program on a distributed-memory environment [1], [2]. However, both existing checkpointing techniques cannot be used for an effective implementation. In case of complete checkpointing that was used in the first version of CAPE the very large generated checkpoints decreased the global performance. Although the incremental checkpointing technique can give smaller checkpoints, it has not enough services for the requirements of the discontinuously and alternatively checkpointing and recovering in CAPE.

As a result, in order to cope with both above requirements, this article presents the new approach and the original optimization we have developed and evaluated with very good first results.

The article is organized as follows: after the related works, Sec. III details our approach to generate discontinuous incremental checkpoints and also provides an evaluation of this approach; Sec. IV develops the data structure in checkpoint files, discussing memory granularity, the different storage format and arithmetics on checkpoints.

II. RELATED WORKS

A. Incremental checkpointing

An incremental checkpoint consists in a file that stores the parts of the memory[3] that have been updated since the beginning of the execution of the program or the last checkpoint. From a high-level point of view, this is performed by setting access rights to read-only to all memory pages in order to force the system to deliver a SIGSEGV signal the next time a page is accessed for writing. Upon reception of the SIGSEGV signal, a copy of the content of the page is stored and access rights to the memory pages are restored to their initial values. When a checkpoint is required, the content of all modified pages are compared to their initial content and the difference is stored in the checkpoint file.

In order to perform these operations, each process that may be checkpointed has to be associated a monitor. Typically, this monitor is in charge of starting the process which checkpoints will be computed, catching the SIGSEGV signals, generating

---

[3]References to the memory are, unless otherwise specified, references to the virtual memory and not the physical memory. In the same way, the paper refers to virtual pages and virtual address spaces and not physical pages and physical address spaces respectively.

[1]viet_hai.ha@it-sudparis.eu
[2]eric.renault@it-sudparis.eu

the checkpoint files and waiting for the termination of the process. However, this is not mandatory as a monitor may be attached to any already running process.

Clearly, the speed of the monitored process decreases due to the execution of the monitor. Even when writing a single byte to a read-only memory region, a set of operations both in the kernel and in the monitor is launched: the kernel issues SIGSEGV; the monitor catches SIGSEGV, sets the access write of the corresponding page to readwrite after reading and saving its initial values. These operations increase the total time to finish the initial writing operation. If parts of the program contain a series of such operations, for example in case of initializing a large memory area from a constant or from values read in a database, this increase may become very important. Column 5, row 3 of Table II on page 5 highlights a case where the execution time is increased to nearly 10 times. The main issue remains on how to avoid this decrease of speed while ensuring the recovering later.

B. CAPE

CAPE [1], [2] stands for Checkpointing Aided Parallel Execution. It consists in modifying a sequential program so that instead of executing each part the one after the other on a single machine, parts are distributed over a set of machines to be executed in parallel. CAPE is not intended to automatically detect which parts of the original code have to be executed in parallel such parts are cited by programmers while using pragma directives of OpenMP. CAPE only aims at transforming a program so that it can be executed in parallel. A typical example is the distribution of a for loop. If the different iterations of a for loop are satisfying the Bernstein's conditions, i.e. any modified memory location is used only in the iteration loop where it is modified, it becomes possible to execute each loop iteration independently on different machines, compute the list of memory locations that have been modified in each iteration loop and include all these modifications inside a single process that will behave as if all iteration loops would have been executed locally.

While using ckpt [7], a complete checkpointer, CAPE has been proved its feasibility. Fig. 2 presents the general template for for loops in form `for ( A; B; C ) D;` when using the C programming language.

In this template, commands `create ( before`$_i$` )` and `diff ( before`$_i$`, after`$_i$`, delta`$_i$` )` aim at extracting `delta`$_i$, the memory areas that have been updated on `host`$_x$, after this host has finished its part of the loop. The use of a complete checkpointer requires an extra time for rendering this value and using an incremental checkpointer can avoid it. Another extra time is caused by the commands `merge ( target, delta )` and `restart ( target )` which serve to inject these `delta` into the memory of the initial host. Ability to directly execute this activity and then allow the process continue running can avoid this extra time. This leads to the requirement of the capacity to take and inject discontinuous incremental checkpoints in programs.

```
parent = create ( original )
if ( ! parent )
        exit
copy ( original, target )
for ( A ; B ; C )
        parent = create ( before_i )
        if ( parent )
                ssh host_x restart ( before_i )
        else
                D
                parent = create ( after_i )
                if ( ! parent )
                        exit
                diff ( before_i, after_i, delta_i )
                merge ( target, delta_i )
                exit
parent = create ( final )
if ( parent )
        diff ( original, final, delta )
        wait_for ( target )
        merge ( target, delta )
        restart ( target )
```

Fig. 1.   Template of CAPE for for loops.

III. DISCONTINUOUS INCREMENTAL CHECKPOINTER IMPLEMENTATION

A. Principle

Based on the incremental checkpointing approach, the main idea for discontinuous incremental checkpointing is to add information to indicate which sections should be checkpointed by the monitor. This information is provided using pragma directives and may be implemented using different mechanisms like signals, breakpoints, etc. Three pragmas have been defined:

- `pragma dickpt start`
- `pragma dickpt stop`
- `pragma dickpt save` $< filename >$

Their behavior is described in Table I. In this table, pages refer to the pages of the virtual address space of the monitored process.

Assume a program consists of segments (A, B, C, D) in which, only B and D need to be checkpointed and two checkpoints are taken in B and one in D. Fig. 2 presents the prototype of the changed program, i.e the directives have been inserted to verify the above requirements.

B. Evaluation of performance

In order to highlight the impact of our new approach, we have developed a new checkpointer named Dickpt and measured its impact on a program computing the successive elements of a Markov Chain, see Fig. 3. Two cases were tested. The first one includes a directive to begin the checkpointing at location 0 (line 10) and takes a checkpoint at location 1 (line 23), while the second one avoids the checkpoint at location 1 and begins checkpointing at location 1. For both cases, one hundred state vectors are computed at location 3 (line 30)

		Monitor	
Pragma	Process	Checkpointing mode	Recovering mode
start	send "start" to the monitor	set all pages to read-only status.	find the next checkpoint: - if found: + inject checkpoint to the monitored process; + set all pages to read-only status; + if it is the last checkpoint: change to checkpointing mode; - else: notice error; stop process.
stop	send "stop" to the monitor.	back all pages to their original status.	back all pages to their original status.
save $< filename >$	send "save" to the monitor.	save the memory locations that have been modified to the current checkpoint; if $< filename >$ contains previous checkpoint: + merge current checkpoint to $< filename >$; else: + save current checkpoint to $< filename >$; set all pages to read-only status.	notice error; stop process.

```
A
# pragma dickpt start
B₁
# pragma dickpt save <filename1>
B₂
# pragma dickpt save <filename1>
# pragma dickpt stop
C
# pragma dickpt start
D
# pragma dickpt save <filename2>
# pragma dickpt stop
```

Fig. 2.   Pseudo-code for discontinuous incremental checkpoints.

```
5 # define LOOP 100
6 # define N     9960
7 # define RAND 10000
8 float M [ N ] [ N ], V [ 2 ] [ N ] ;
9 int main ( int argc, char * argv [ ] ){
10   //location 0
11   float sum, val ;
12   int i, j, k, l ;
13   for ( i = 0 ; i < N ; i ++ ) {
14     for ( sum = j = 0 ; j < N ; j ++, sum += val )
15       M [ i ] [ j ] = val = rand ( ) % RAND ;
16     for ( j = 0 ; j < N ; j ++ )
17       M [ i ] [ j ] /= sum ;
18   }
19   for ( sum = i = 0 ; i < N ; i ++, sum += val )
20     V [ 0 ] [ i ] = val = rand ( ) % RAND ;
21   for ( i = 0 ; i < N ; i ++ )
22     V [ 0 ] [ i ] /= sum ;
23   //location 1
24   for ( k = l = 0 ; l < LOOP ; l ++, k = 1 - k ) {
25     for ( i = 0 ; i < N ; i ++ ) {
26       V [ 1 - k ] [ i ] = 0. ;
27       for ( j = 0 ; j < N ; j ++ )
28         V [ 1 - k ] [ i ] += ( V [ k ] [ j ] * M [ j ] [ i ] ) ;
29     }
30     //location 2
31   }
32   return 0 ;
33 }
```

Fig. 3.   Program computing the successive elements of a Markov Chain.

and one checkpoint is generated after each computation. The testbed was composed of an Intel Core2 Duo E8400 running at 3 GHz with 3 GB RAM and operated by Ubuntu 9.10 based on Linux kernel 2.6.31-21-generic. Table II presents the performance evaluation for four vector sizes (N equals to 3320, 6640, 9960 and 13280 elements respectively). For each vector size, performance are measured 30 times (mean values are provided in the table). In order to avoid the pollution of disk effects on measurements, all data (the virtual address space of processes, checkpoints, etc.) are resident in RAM.

The first section of Table II presents the size of checkpoints at location 1 (i.e. just after the initialization - as the case of normal incremental checkpointer) and at location 2 (i.e. just after the computation of a new vector). This difference is the size of the transition matrix which is initialized at the beginning of the program. These data show how much disk space can be saved while abandoning the checkpoint at location 1 and, in the same way, how faster the checkpoint can be transfered over the network if necessary.

The second section of the table shows the time required to run the program without saving checkpoints, while saving all checkpoints and while saving location 2 checkpoints only. It

highlights the fact that the overhead involved by the generation of location 2 checkpoints is very light (between 1.5% and 3.3% for 100 checkpoints, i.e. between 0.01% and 0.03% per checkpoint) compared to the overhead involved by the generation of both location 1 and location 2 checkpoints (the total execution time is multiplied by 8.5).

The third section of table provides the execution time to run the process restarting from loop iteration number 50. Two cases are envisaged: the first one uses all checkpoints, i.e. the program is restarted, suspended at the beginning of function main, all checkpoints are injected in the process and the execution resumes at loop iteration 50; the second one uses location 2 checkpoints only, i.e. the program is restarted, suspended after the initialization step, all location

2 checkpoints are injected in the process and the execution resumes at loop iteration 50.

The performance measurements show that avoiding location 1 checkpoints is always beneficial.

C. Advantages and drawbacks

While comparing with the normal incremental checkpointing technique, our new approach has the strong and weak points below:

- Performance: increase the speed of the program in both periods of checkpointing and recovering; decrease strongly the size of checkpoints.
- Flexibility: allows to select the segments to be checkpointed in programs. The case of normal incremental checkpointing is obtained by setting a `pragma dickpt start` and a `pragma dickpt stop` as the first and the last instruction respectively in the checkpointed program.
- More specifically for the case of CAPE: it can directly extract and allows to inject the `delta`, thus increase the performance of this paradigm.
- Change of the source code: in the role of a checkpointer, it is the most important drawback. Users have to insert the directives to indicate the regions being checkpointed. However, when used in CAPE, this insertion is done by the CAPE's compiler. As a result, this drawback has no impact for CAPE's users.
- Fragmentation of checkpoints: checkpoints which are not taken in a single block (surrounded by a pair of `pragma dickpt start` and `pragma dickpt stop`) can not be merged to an unique checkpoint. So, many files are needed to contain the checkpoints of different checkpointing blocks. This drawback is important when checkpoints reference the same memory area.

IV. CHECKPOINT STRUCTURE OPTIMIZATION

The structure of a complete checkpoint is usually quite straightforward. After some very specific data like the content of registers and the size of the memory, the rest of a complete checkpoint is usually composed of the content of all the memory pages the one after the other one.

In the case of an incremental checkpoint, several cases have to be envisaged. All solutions are storing the content of registers. However, regarding the memory updates, the best solution really depends upon the granularity of data, which ranges from one byte to one page with the most interesting case at one word.

Memory granularity

There are two main drawbacks when the granularity is the page. The first one is that a complete page must be saved even though a single byte in the page has been modified, which is not really memory efficient. Considering the size of today disks, this may not be a problem unless a very large number of checkpoints have to be generated. The problem may have a more important impact if for examples these checkpoints have to be sent over the network, especially with a limited bandwidth. The second main drawback is that it provides no information on which bytes in the page have been modified effectively. The latter drawback definitely forbids any merge operation of successive incremental checkpoints.

Setting the granularity of the checkpoint to a single byte solves the memory inefficiency problem of the page granularity. However, it leads to other subtle problems, like for example the reference to memory locations that do not exist in the virtual address space of the process. Let $< a, b, c, d >$ be four bytes stored at a memory location and representing a pointer in memory. After a first checkpoint, this memory location may contain $< a, b', c, d >$. After a second checkpoint, the same memory location may contain $< a, b, c', d >$. If, for any reasons, it is required to merge the two checkpoints (and this is typically the case with CAPE), the result becomes $< a, b', c', d >$ which may not be part of the virtual address space of the process.

Setting the granularity of the checkpoint to a word (i.e. four bytes) is the best compromise as it solves the problem of memory space efficiency and does not introduce any pointer problem as described above. This solution is not the perfect solution. However, problems involved by setting the granularity of the checkpoint to a word has no significant impact on the execution of the program.

Finally, one can note that setting the granularity of the checkpoint to the entire virtual address space turns an incremental checkpointer into a complete checkpointer.

Incremental checkpoint content

Apart from the specific values also stored in complete checkpoints, an incremental checkpoint should be composed of the list of memory locations that have been modified since the beginning of the execution of the program, or since the previous checkpoint, and the last value for each of these specific memory locations. The simplest structure to store such a list is to save the one after the other one both the addresses and their associated value. However, since the spatial locality of data in most programs implies that a modification at a memory location increases the probability for adjacent memory locations to be modified, this way of storing data is not necessarily efficient.

Thus, in order to take advantage of the spatial locality of updates and therefore reduce the size of checkpoints, several alternative methods for storing memory updates have been identified:

- *Single data*. This case occurs when a single memory location has been updated. In this case, the only information to store are the basic address of the memory location and the content at the memory location. Data to store all information into the checkpoint are:

$$< \text{addr}, \text{value} >$$

- *Several successive data*. This case occurs when more than one consecutive memory locations have been updated. For example, this is encountered when the content of an array has been modified. The best way to store all the information in this case is:

TABLE II
PERFORMANCE EVALUATION.

	Matrix size			
	3320	6640	9960	13280
Checkpoint size (in MB)				
... at location 1	42.192	168.741	379.648	674.912
... at location 2	0.013	0.026	0.038	0.051
Total execution time (in seconds)				
... without generating checkpoints	11.34	41.30	108.27	168.69
... generating all checkpoints	13.10	58.14	393.60	1433.72
... generating location 2 checkpoints only	11.71	42.16	109.96	171.70
Execution time restarting after loop iteration #50 (in seconds)				
... using location 1 and location 2 checkpoints	6.41	23.14	59.56	93.91
... using location 2 checkpoints only	6.09	21.84	56.64	88.71

$$< \text{addr}, \text{size}, [\text{value}...] >$$

- *Many data*. This occurs when lots of non-successive memory locations have been updated on a single page. In this case, instead of storing a large number of Single data and Several successive data elements, it is more efficient to store the address of the page, the list of memory locations on the page that have been modified and for each modified memory location the associated value. The efficiency of this solution resides in the mapping, i.e. the list of memory locations on the page. As this is a binary information for each data in the page, it can be represented using a single bit per memory location. For example, for a 4-kB page, the size of the map is 1024 bits (or 128 bytes) with a granularity set a word.

$$< \text{addr}, \text{map}, [\text{value}...] >$$

- *Entire page*. This occurs when all memory locations on a memory page have been modified. This case is quite common when a new page is added to the virtual address space of a process. The best way to store the complete content of a page is:

$$< \text{addr}, [\text{value}...] >$$

No size need to be provided in this case as it is implicitly known.

Table III compares the amount of memory needed to store updated data for all cases presented above. The size of a memory page is assumed to be 4 kB. Let a *chunk* be a set of contiguous memory locations that have been updated. Let c be the number of chunks in a memory page, let s_i be the number of elements in chunk i and let u be the number of updates in the memory page. By definition, $\sum_{i=1}^{c} s_i = u$.

TABLE III
AMOUNT OF MEMORY TO STORE UPDATES.

Method	Amount of memory	
	for a single chunk	for a page
Single data (SD)	8	$8 \times u$
Several successive data (SSD)	$8 + 4 \times s$	$8 \times c + 4 \times u$
Many data (MD)	$132 + 4 \times u$	$132 + 4 \times u$
Entire page (EP)	4100	4100

Fig. 4 shows a comparison of the amount of memory needed to store all updates in a 4-kB page as a function of the number of updated memory locations in the page. SD, MD and EP only depend upon the number of updated memory locations while SSD also depends on the distribution of the updated memory locations. As a result, Fig. 4 shows both the best case (SSD_{min}) that is when all updated memory locations are in a single chunk, and the worst case (SSD_{max}) that is the case when updated memory locations are distributed in the configuration that requires the maximum number of chunks. For 4-kB memory pages and 4-byte words, this maximum is given by:

$$\begin{cases} \lfloor u/2 \rfloor & \text{if } 0 < u \leq 682 \\ 1024 - u & \text{if } 682 < u \leq 1024 \end{cases}$$

One can note that when two successive memory locations have to be stored, the amount of memory needed to store the information for both Single data and Several successive data cases is the same.


Fig. 4. Amount of memory to store updates.

The most efficient solution, i.e. the one that reduces the most the memory usage, is identified this way. For each page, first the Many data representation is built. It requires at most 4228 bytes; second, a combination of Single data and Several successive data methods is built, having Single data chosen for isolated data and Several successive data chosen when at least two consecutive memory locations have been updated; third,

the shortest representation between both computed is stored. Note that the Entire page method is left to the storage of new pages.


Fig. 5.   Trade-off between SSD and MD.

From the expressions provided in Table III one can demonstrate that the Several successive data method is always the most interesting solution when the number of updates is smaller than 34. Then, there is a trade-off between the Several successive data method and the Many data method that depends on the number of chunks. Fig. 5 is a magnification of Fig. 4 for a number of updates in the range from 0 to 50.

Identifying the method

Considering that more than one method is used to store memory updates, it is important to identify which one was used when restoring the content of the checkpoint. A simple solution would have consisted in adding an extra integer or even a character before any data description or set of data description. However, in order to keep the size of the checkpoint as small as possible, it has been decided to add no extra byte to the checkpoint.

Instead, considering that all methods require an address as the first field and that these addresses are necessarily aligned on a boundary of a word, i.e. these addresses are necessarily a multiple of 4 or the last two digits of their binary representation are necessarily 00, it is possible to use this "free" space to store which method was used to store the data. In our current implementation, 00 is associated to Single data, 01 to Several successive data, 10 to Many data and 11 to Entire page. When restoring the content of a checkpoint, these two bits are reset to 00 after the storage method has been identified and before the address is effectively used.

V. Conclusion and future works

For the development of our distributed implementation of OpenMP using the CAPE paradigm, an efficient incremental checkpointer is required. This article presented our new approach of discontinuous incremental checkpointer. The initial implementation of this approach, Dickpt, has proved its efficiency in terms of size - the decrease of amount of memory required to store the checkpoints - and time to

generate the checkpoints, to restart from a checkpoint or to send the checkpoint over the network. This approach led to very lightweight checkpoints, the size being in the order of very few kB while the size of the virtual address space of the process is in the order of tens or hundreds of MB. However, by the significantly changes due to the possibility and the nature of the checkpointer, new paradigms may be developed for CAPE and we are investigating in this way.

Apart from processor registers, the current implementation of our incremental checkpointer does not save system information about the process, like open file descriptors, sockets, POSIX semaphore or shared memory, etc. Somehow, this is not a real drawback as if such declarations are set outside of checkpointed segments, they could be re-executed in case of restoration. However, in the near future, these system data will be included and the overhead involved by the inclusion of each of them will be studied. In fine, we would like this incremental checkpointer to be completely customizable so that users would be able to store only relevant information.

References

[1] Eric Renault. *Parallelization of For Loops Using Checkpointing Techniques*. Proceedings of the 2005 International Conference on Parallel Processing Workshops, Oslo, Norway, pp 313-319, June 2005.
[2] Eric Renault. *Distributed Implementation of OpenMP Based on Checkpointing Aided Parallel Execution*. International Workshop on OpenMP (IWOMP), Beijing, China, LNCS 4935, pp 183-193, June 2007.
[3] L. Mereuta and Eric Renault. *Checkpointing Aided Parallel Execution Model and Analysis*. High Performance Computation Conference (HPCC), Houston, TX, LNCS 4782, pp 707-717, September 2007.
[4] http://www.research.rutgers.edu/ edpin/epckpt/
[5] J. Ansel, K. Arya and G. Cooperman. *DMTCP: Transparent Checkpointing for Cluster Computations and the Desktop*. 32rd IEEE International Parallel and Distributed Processing Symposium (IPDPS09), Rome, Italy, May 2009.
[6] http://cryopid.berlios.de/
[7] http://pages.cs.wisc.edu/ zandy/ckpt/
[8] http://hpc.pnl.gov/sft/tick.html
[9] S. YI, J. Heo, Y. Cho, J. Hong, J. Choi and G. Jeon. *Ickpt: An Efficient Incremental Checkpointing Using Page Writing Fault - Focusing on the Implementation in Linux Kernel*. Proceedings of the ISCA 19th International Conference on Computers and Their Applications (CATA04), Seattle, WA, pp. 209-212, March 2004.
[10] J.S. Plank, M. Beck, G. Kingsley and K. Li. Libckpt: Transparent Checkpointing under Unix. Proceedings of the Usenix Winter 1995 Technical Conference, New Orleans, LA, pp. 213-223, January 1995.
[11] K.A. Iskra, F. van der Linden, Z.W. Hendrikse, B.J. Overeinder, G.D.van Albada and P.M.A. Sloot. *Sloot: The implementation of Dynamite an environment for migrating PVM tasks*. Operating System Review, 34:3(40-55), ACM SIGOPS, July 2000.
[12] http://checkpointing.psnc.pl/Progress/psncCR/
[13] O.O. Sudako, I.S. Meshcheriakov, Y.V. Taras ShevchentoA. *Transparent Checkpointing System for Linux Clusters*. 4th IEEE worksho on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS). 2007.
[14] Y.-M. Wang, Y. Huang, K.-P. Vo, P.-Y. Chung, and C. Kintala. *Checkpointing and its applications*. 25th Annual Intl symposium on Fault-Tolerant Computing, pp.22-30, Oct. 1995.
[15] http://openmp.org/wp/