Detection of Silent Data Corruptions in Smoothed Particle Hydrodynamics Simulations

Aurélien Cavelan University of Basel, Switzerland aurelien.cavelan@unibas.ch

Rubén M. Cabezón University of Basel, Switzerland ruben.cabezon@unibas.ch

Florina M. Ciorba University of Basel, Switzerland florina.ciorba@unibas.ch

April 24, 2019

Abstract

Silent data corruptions (SDCs) hinder the correctness of long-running scientific applications on large scale computing systems. Selective particle replication (SPR) is proposed herein as the first particle-based replication method for detecting SDCs in Smoothed particle hydrodynamics (SPH) simulations. SPH is a mesh-free Lagrangian method commonly used to perform hydrodynamical simulations in astrophysics and computational fluid dynamics. SPH performs interpolation of physical properties over neighboring discretization points (called SPH particles) that dynamically adapt their distribution to the mass density field of the fluid. When a fault (e.g., a bit-flip) strikes the computation or the data associated with a particle, the resulting error is silently propagated to all nearest neighbors through such interpolation steps. SPR replicates the computation and data of a few carefully selected SPH particles. SDCs are detected when the data of a particle differs, due to corruption, from its replicated counterpart. SPR is able to detect many DRAM SDCs as they propagate by ensuring that all particles have at least one neighbor that is replicated. The detection capabilities of SPR were assessed through a set of error-injection and detection experiments and the overhead of SPR was evaluated via a set of strong-scaling experiments conducted on an HPC system. The results show that SPR achieves detection rates of 91-99.9%, no false-positives, at an overhead of 1-10%.

1 Introduction

The frequency of faults, errors and failures in large-scale systems increases proportionally to the number of components. As such, reliability has been identified as a major challenge for Exascale computing [14,15,43]. Trustworthy computing aims at guaranteeing the correctness of the results of long-running computations on large-scale supercomputers. In particular, silent errors, or *silent data corruptions* (SDCs), represent a major threat to the correctness of the results of such calculations. SDCs typically manifest when one or more bits are flipped in the memory of the system. There are several causes for these errors, such as packaging pollution or cosmic rays, among others [4, 39, 51].

The standard solution for protecting the memory subsystem against data corruption is the use of error correcting codes (ECC). However, these mechanisms are not able to detect and correct all errors. Indeed, recent studies suggest that such mechanisms may not prevent data corruptions from occurring at extreme scales [4,33,37], and in particular in DRAM devices [45]. This work expands on detecting DRAM errors.

Triple modular $redundancy^1$ (TMR) [35] is the most general and non-intrusive approach to detect and correct SDCs in scientific applications. SDCs detectors are characterized by their precision and their recall. The precision is the ratio of true errors detected (in contrast to false-positives) over all errors detected, and the recall is simply the ratio of true errors detected over all true errors that occurred during the execution. A high precision and a high recall indicate both few false-positives and a good detection rate, respectively. In general, detectors that either employ full replication of entire applications [27] or selective replication of parts of an application [10] offer the highest precision and recall. However, they are often prohibitively expensive in terms of additional required computing resources and time. Therefore, many application-specific detectors have been proposed as alternatives to redundancy to lower the cost of error detection. Specifically, data analytics-based techniques detect outliers by relying on application-specific properties, such as spatial and/or temporal data smoothness. Interpolation-based detectors employ techniques such as time series prediction and spatial multivariate interpolation [2,5,6] to interpolate the next value of a data-point, and offer broad error detection coverage at low cost, but typically have a lower precision compared to full-replication approaches. A further review of relevant error detection methods is presented in Section 4.

Smoothed particle hydrodynamics (SPH) is a meshless Lagrangian method commonly used for performing hydrodynamical simulations. Initially devised in the late '70s [31, 34], this technique underwent sustained development [13, 28,41,44] and is, nowadays, commonly used in many fields including computational fluid dynamics, plasma physics, solid mechanics, and astrophysics. Its inherent good conservation properties and its adaptability to distorted geometries make SPH a common choice to simulate highly-dynamic three-dimensional scenarios. The SPH technique discretizes a fluid in a series of interpolation points (called SPH particles) whose distribution follows the mass density of the fluid, and their evolution relies then on a weighted interpolation over close neighboring particles. The cost of such a highly-adaptive unstructured fluid discretization is the *non-uniformity of data*, in terms of the number and distribution of neighbors that a particle might have at a specific location and time. This renders most data analytics-based error detectors less accurate than redundancy-based detectors. Moreover, SPH simulations are inherently robust: it is likely that SDCs will not crash the execution of an SPH application, thereby not alerting the user.

¹In this work, *redundancy* and *replication* are used interchangeably.

Selective particle replication (SPR) is proposed in this work. SPR replicates the computation and data of a few carefully selected SPH particles, and SDCs are detected when the data of a particle differs from the data of its replica or viceversa.

It is necessary to understand how SDCs propagate in SPH simulations before we clarify how to select which particles need to be replicated. We assume that an SDC initially strikes the data of a single particle (e.g., a bit-flip in the position, density, mass, temperature, or in any other data field associated with that particle). The subsequent SPH interpolation steps will proceed with erroneous data and propagate the error to all neighbors of that particle, i.e., typically affecting ~ 100 other particles. However, having just one redundant neighbor would be enough to detect the difference (no matter how small) in the computed result due to the original SDC. Therefore, by ensuring that *all particles have at least one redundant neighbor*, SPR can detect most SDCs (as they propagate). We detail the entire process and its limitations in Section 2.

We conduct a series of experiments by incorporating SPR in a production SPH code and running two test simulations commonly used in astrophysics. First, we perform a set of error-injection and detection experiments to assess the detection capabilities of SPR. Next, we conduct a set of performance-evaluation experiments, without injecting errors, to assess the overhead of SPR for the two test simulations on the Piz Daint supercomputer². The outcome shows that SPR is particularly suited for detecting SDCs in SPH simulations, detecting 91 – 99.9% of the injected errors, for a typical execution overhead between 1 - 10%. In addition, SPR raises no false-positives, which is a significant advantage over unprecise detectors.

This work makes the following contributions: (1) Presents a simple, yet effective method (SPR) for selecting particles to be replicated and detects errors with bit-level precision; (2) Experimentally quantifies the recall and precision of SPR through an error-injection and detection campaign in a production scientific application; and (3) Assesses the scalability and overhead of the proposed SPR approach through real experiments on an HPC system.

The remainder of this work is structured as follows. We describe in detail and illustrate the proposed approach for silent error detection via *selective particle replication* in Section 2. We include the experimental setup and results in Section 3 along with their evaluation and a discussion of the findings. Section 4 surveys the relevant literature related to silent errors detection. Finally, Section 5 outlines the conclusion of this work together with insights into future work extensions.

2 Selective Particle Replication

Selective particle replication (SPR) consists of three steps, namely selection, replication, and detection. Section 2.1 describes an algorithm to select the particles that will be replicated, Section 2.2 covers the different aspects and challenges of particle replication in SPH, and Section 2.3 presents the error-detection algorithm as well as its limitations.

²https://www.cscs.ch/computers/piz-daint/

2.1 Selection



Figure 1: SPR step 1: selection of particles to replicate. The figure shows the distribution of 32 particles on three processes $\{q - 1, q, q + 1\}$ computing on particles from different sections of the SPH domain. Selected particles on each process form a Maximal Independent Set (MIS) such that all particles have at least one neighbor that is selected for replication in another process. Edges show the connection between the selected particles and their neighbors.

In parallel SPH implementations [38], the particle domain is distributed among the processes. Each process is given a list of particles and their associated data. The first step in SPR is to select which particles to replicate to ensure that all particles in a process have at least one neighbor that is replicated in another process, as illustrated in Figure 1.

Given that the number of computations grows linearly with the number of particles selected on each process, we want to replicate as few particles as possible overall. The domain composed of the particles and their neighbors can be compared to a graph where vertices and edges represent the particles connected to their neighbors.

Particles must be selected such that any given particle in a process either belongs to the set of particles selected for replication, or has a selected particle as neighbor, which translates to finding a *Minimum Dominating Set* (MDS). However, this is a classical NP-Hard problem [29] and there is no known optimal polynomial time algorithm for finding such a minimum set. One can alternatively solve the simpler *Maximal Independent Set* (MIS) problem, where a MIS is also a Dominating Set (albeit not minimum). As opposed to a dominating set, a MIS ensures that no selected vertices (respectively particles) in the graph are adjacent, which leads to fewer vertices being selected, and it is possible to find many MISs in a graph.

The present work proposes the use of a greedy algorithm to find a MIS, listed in Algorithm 1. Given the set P_q of particles assigned to process q for computation, Algorithm 1 returns the set S_q of particles selected for replication in another process. A possible solution to the problem is presented in Figure 1 for a small example with 32 particles distributed on three processes, highlighting the selected particles in each process.

Complexity. Algorithm 1 requires, in the worst case, to iterate through the n_q particles of process q and for each particle to visit at most ng_{max} neighbors,

Algorithm 1 Particle Selection

procedure SELECT-PARTICLE (P_q)
Initialize S_q to an empty set
while P_q is not empty do
Choose a particle $p \in P_q$
Add p to S_q
Remove from P the particle p and all its neighbors
end while
$return S_q$
end procedure

where ng_{max} denotes the maximum number of neighbors for any given particle in the simulation. Therefore, the running time of this algorithm is $\mathcal{O}(n_q n g_{max})$. **Overall asymptotic cost.** The running time of SPH applications is dominated by large interpolation kernels scaling as $O(n_q n g_{max})$. Comparatively, our approach only requires $\mathcal{O}(n_r n g_{max})$ additional computations, where n_r is the maximum number of replicated particles per node. A lower bound on n_r is $\frac{1}{ng_{max}}$, when only 1 particle is selected to dominate all of its neighbors. Experiments show n_r to be in the range of 1 - 10% based on our MIS approach. Consequently, the increment in memory cost and communication time is also very small 1 - 10% (see Section 3).

Note that the maximum number of neighbors, ng_{max} , is often a user-defined parameter in the order of 10^2 particles for 3D simulations.³ As the simulated scenarios grow in complexity and require increased resolution, the overall number of particles in SPH simulations increases, as well as the number of neighbors. Nowadays, 3D SPH simulations require between $10^5 - 10^{12}$ particles, with $10^2 - 5 \cdot 10^2$ neighbors per particle.

2.2 Replication

The second SPR step is to copy the data and to replicate the execution of the previously selected particles onto a target process. Each process must send a copy of all its selected particles and their neighbors to its assigned target process. As the selected particles are in fact a MIS, each process copies of the data of *all* its particles on to its assigned target process, as shown in Figure 2. However, only the selected particles are computed on the target process.

Figure 2 shows the distribution of 32 particles onto three processes $\{q - 1, q, q+1\}$. In this configuration, each process initially sends a copy of the data associated with all its particles to the next process id. Note that both computation and data are replicated onto a target process for the selected particles, while only data is replicated for all the neighbors of a selected particle. This data is needed to update the selected particles during SPH neighbor interpolation steps.

To avoid the risk of using corrupted data, detection must be performed before any inter-process communications occur. As long as no data is communicated to another process, the error is guaranteed to be contained within the faulty process. Since in SPH, processes must exchange data after every interpolation

 $^{^{3}}$ Given a target number of neighbors, the simulation will try to reach this number of neighbors for each particle and this influences the resulting smoothing length.



Figure 2: SPR step 2: replication of selected particles onto the target process. Each process sends a copy of the data associated with all its particles to the next process id. Note that all particles are copied, but only selected particles (orange, green, blue) are computed upon the target process, using the data from their replicated neighbors (light gray).

step to exchange neighbors that are at the edge of their respective computational domain, SPR detection must also be done right after each interpolation step. In addition, each process must update the copy of the data of all its particles on the target process, so that the selected replicated particles use the latest version of their neighbors data.

The general workflow is described in Algorithm 2 and additional SPR steps are shown in blue. Depending on the scenario they simulate or the research field in which they are used, SPH codes can greatly vary in terms of implementation and physical processes they include. Nevertheless, many SPH codes apply the same general underlying workflow. This allows us to generalize the use of the SPR method proposed here, not only to a single specific SPH code but to the vast majority of SPH codes.

2.3 Detection

Error-detection is done by comparing the data of the selected particles against their the data of their replica. Detection must be performed after every computational workflow step, before any data are communicated to other processes, as illustrated in Algorithm 2 (in blue).

Given that an original particle and its replica represent the same particle, their data are expected to remain identical at any point during the simulation.

3

Initialization while Target simulated time is not reached do 1. Build tree 1.1 Detection 1.2 Update of neighboring particles 2. Find neighbors 2.1 Detection 2.2 Select particles for replication 2.3 Update of neighboring particles 3. Execute SPH neighboring interpolation kernels 3.1 Detection 3.2 Update of neighboring particles 4. Find new time-step 4.1 Detection 4.2 Update of neighboring particles 5. Update velocity and position 5.1 Detection 5.2 Update of neighboring particles 6. (Optional) Compute self-gravity 6.1 Detection 6.2 Update of neighboring particles end while

Moreover, the data of the original particle can be compared against the data of its replica with *bit-level resolution*. If an error propagates to either the original particle or its replica, even by causing a change in single bit, it will be detected via a simple comparison.

Figure 3 shows how one of the selected particle on process q becomes corrupted after using data from a corrupted neighbor during an interpolation step. Since the newly computed data on process q differs from the data of its replica on process q + 1, the error is detected via a simple comparison.

The following algorithm is proposed to detect errors using replicated particles selected a priori in Step 1. We assume that replicated particles are always computed upon the next process in a round robin fashion, i.e., the list of replicated particles selected on process q with Algorithm 1 is sent to process q+1, so that process q+1 computes the selected replicated particles from process q, and so on. Therefore, for detection purposes, process q sends the data corresponding to replicated particles from process q-1 back to process q-1, and in return, it receives the data corresponding to its own replicated particles from process q-1 and let $data_q(S_{q-1})$ denote the data associated with replicated particles selected on process q and used for computation on process q. Similarly, let $data_{q+1}(S_q)$ denote the data associated with replicated selected on process q and used for computation on process q+1. The detection algorithm 3.

For a given process q, Algorithm 3 iterates through the list of replicated particles that have been selected on process q and computed on process q + 1.



Figure 3: SPR step 3: error detection. An error (e.g. a bit flip) has corrupted the data associated with a particle (in red). When the selected particle (in green) that is a neighbor of the corrupted particle is updated via interpolation over its neighbors, the error is propagated and the data of the selected particle becomes corrupted. SPR detects the effects of this silent error from the corrupted particle by comparing the data of the selected particles on q with their replica on the target process q + 1.

For each selected particle, it compares the data obtained on process q + 1, corresponding to the replica, with the data obtained on process q, corresponding to the original particle. If the data differ, an error has been detected.

2.4 Limitations

The SPR detection capabilities are limited to errors that propagate. SPH codes typically simulate several complex physical phenomena in the same application. Some particle datasets, such as the positions of the particles, are used very often during interpolation throughout the simulation. The data is read and used multiple times in different interpolation steps. An error in such datasets is more likely to propagate, and therefore to be detected, while errors in datasets than are used less often, such as particle particle densities, are less likely to propagate and may not always be detected.

Note that this approach is limited to deterministic calculations and that round-off errors are expected and need to be accounted for. This is typically achieved by allowing k lower order bits to differ. The SPH method is entirely

Algorithm 3 Error-detection algorithm for process q

```
procedure ERROR-DETECTION(S_q, S_{q-1})
Set error to false
Send data_q(S_{q-1}) to process q-1
Receive data_{q+1}(S_q) from process q+1
for particle p in S_q do
if data_{q+1}(p) differs from data_q(p) then
Set error to true
end if
end for
return error
end procedure
```

deterministic, and in the experiments performed in this work, no round-off errors or truncated values were encountered.

2.5 Error Correction

Even though the main focus of this work is on the detection of SDCs, it is possible to combine the proposed detector with other error-detection and correction methods. Checkpointing with rollback recovery [17, 23] is the de-facto general-purpose recovery technique in high-performance computing. Finding the optimal checkpointing interval [7, 19, 21, 49] or the optimal recovery method for SPH codes is beyond the scope of this paper.

Because the SPR error detector is used in every simulation time-step, it is possible to safely checkpoint the state of the simulation if no error was detected. Then, whenever a new error is detected, the simulation can simply rollback to the last correct checkpoint and re-execute from there. For completeness, this work implements checkpointing using the Fault-Tolerance Interface (FTI) [3].

3 Experimental Evaluation

In this section, we describe the experiments conducted with two test simulations commonly used in astrophysics. The goal of these experiments is two-fold: (1) experimentally assess the detection capabilities of SPR through an error injection campaign; and (2) assess the performance of the proposed approach at scale.

3.1 Experimental Setup

SPR implementation into a production SPH code SPR has been incorporated in SPHYNX⁴ [12], an SPH code with focus on astrophysical simulations. The incorporation of SPR into SPHYNX comprises the implementation of Algorithms 1, 2, and 3 presented in Section 2. SPHYNX includes state-of-the-art SPH methods that allow it to address subsonic hydrodynamical instabilities and strong shocks, which are ubiquitous in astrophysical scenarios. SPHYNX is

⁴Freely available at https://astro.physik.unibas.ch/sphynx

implemented in Fortran and uses the message-passing interface (MPI) for interprocess communication and the open multi-processing interface (OpenMP) for intra-process/inter-threads data sharing. Note that compared to the original SPHYNX code, which counts $\sim 30,000$ lines of code, less than 300 additional lines of code were added by SPR as a module for SPHYNX. In the following, the term SPHYNX denotes the original code, and SPHYNX+SPR denotes the version of the code extended with SPR. SPHYNX was used to run two different test simulations: (1) the Evrard collapse (EC) [25], which studies the gravitational collapse of a gaseous cloud and is a common test to evaluate the correctness of the coupling of hydrodynamics and self-gravity; and (2) the wind-bubble interaction test (WB) [1], which studies the interaction of a supersonic wind with a high-density colder bubble in pressure equilibrium.

Choice of number of neighbors The number of neighbors per particle must cover the area of influence of the interpolation kernel in a reasonably homogeneous way, and is, therefore, linked to the type of SPH kernel used in the simulation. In general, 10^2 neighbors is a very common value found in the bibliography, for frequently used SPH kernels in 3D simulations. Unlike the kernel used in SPHYNX (Sinc kernel [13]), other SPH kernels, e.g., Wendland kernel [20], have a much longer radius and, as a consequence, must use larger amounts of neighbors (e.g., 5×10^2). These choices are made after an empirical iterative process in the SPH history and are proven to be adequate to balance accuracy and computational cost, while suppressing numerical instabilities, like particle pairing.

Experiments For each test simulation, different parameters were passed to SPHYNX: a file containing the initial conditions, the corresponding number of particles, and other parameters specific to each test. During the time this work was performed, the initial conditions for only three tests cases were available: an EC test simulation with 65,536 particles that was used to experiment with error-injection and detection, and two larger EC and WB test simulations with 1,000,000 and 3,157,385 particles, respectively, that were used to assess the scalability of SPR. Generating initial conditions for different numbers of particles is a *non-trivial* process. Therefore, this work employs a set of strong-scaling experiments to assess the performance at scale for the proposed SPR method. Table 1 summarizes the different experiments and the corresponding parameters used. Error-injection and error-detection experiments were conducted on a small scale system with 20 nodes named *miniHPC*, while strong-scaling experiments were done the *Piz Daint* supercomputer with up to 256 nodes.

3.2 Error-Injection and Error-Detection Experiments

In this section, an error-injection and detection campaign was performed using the ECsg test simulation described in Table 1. The goal was to experimentally assess the detection capabilities of SPR. Specifically, the focus was on measuring the recall and the precision of the detector, where the precision is defined as $100 \times \left(1 - \frac{\# \text{false-positives}}{\# \text{errors detected}}\right)$, and the recall is defined as $100 \times \frac{\# \text{errors detected}}{\# \text{errors detected}}$.

 $[\]frac{1}{\text{#errors detected} + \text{#errors undetected}}$

Experiment target	Error-injection	Strong-scaling	
Test simulation	ECsg	EC	WB
#Processes	16	4-256	4-256
#Threads/process	20	12	12
#Particles	65,536	1,000,000	3, 157, 385
Self-gravity	Yes	Yes	No
HPC system	miniHPC	Piz Daint	
#Neighbors	100		

Table 1: Design of target experiments

A golden set of results was first created by running two time-steps of the application in an error-free environment. To inject errors, an additional thread was created in each process to inject single bit-flips into the data of the application during its execution. The timing of the error to be injected is *critical*: if it is injected too early in a time-step, the error may be overwritten and masked, while if it is injected too late (after propagation occurred), it may not be detected by SPR.

The simulation repeatedly executed the same two time-steps as in the golden set: a single bit-flip was injected at a random time and a random data location within the first time-step, and results were collected at the end of the second time-step. The injected error was *masked* if all datasets were identical to the ones in the golden set, *detected* if SPR raised a flag, and *undetected* otherwise. *False-positive* errors are also counted, which would have been observed if SPR had raised a flag while there was no error. However, *no false-positives were observed in the experiments*, which means that SPR has an experimentally measured precision of 100%.



Figure 4: Fraction of errors detected, undetected, and masked for five particle datasets, where (a) only accounts for significant errors (in the sign, exponent, or first 5 digits), while (b) accounts for all errors (in any of the 16 decimal digits). 25,000 errors were injected over the execution of 50,000 simulation time-steps.

To quantify the recall, experiments were performed of the most five critical particle datasets, namely position, mass, internal energy, velocity and density. Each one of these datasets holds as many 64-bits double-precision numbers as there are particles in the simulation, and up to three times more for multidimensional datasets, such as positions and velocities. Following the IEEE 754 number representation, bit 64 encodes the sign, bits 63 to 53 (included) encode the exponent, and bits 53 to 1 encode the fraction or mantissa, which corresponds to approximately 16 decimal digits of precision. For each of these datasets, 5,000 bit-flips were injected in random processes, particles, and bit positions, for a total of 25,000 errors injected over the execution of 50,000 simulation time-steps.

The results of the error-injection and detection campaign are presented in Figure 4. Note that certain SDCs fall into the significant errors category. These errors correspond to an absolute difference greater than 10^{-6} between the corrupted and the golden data. Most detection techniques focus only on significant errors and will fail to detect less significant errors (of which many of them may be impactful). In fact, only a comparison with the actual correct data can reveal them, which SPR is able to perform.

The differences in the number of detected, masked, and undetected errors for different datasets are due to several factors. Certain particle datasets, such as positions, are critical to the SPH method. They are used in almost every computational step of the workflow, which greatly increases the probability that an error is propagated, and therefore detected by SPR. Other particle datasets, such as density, can be recomputed from other datasets, and are less frequently used in the computational workflow. Therefore, the probability that an error is masked is higher, and there are fewer opportunities for SPR to detect it.



(a) Average execution time per time-step (\overline{B}) Average execution time per time-step (WB)

Figure 5: Average time needed to complete one time-step for the EC test with the resilient code (SPHYNX+SPR) and the original version (SPHYNX) for the EC test (a) and the WB test (b).

3.3 Strong Scaling Experiments

Overall, the measured recall was always greater than 0.91, and as high as 0.999 for significant errors. The recall is only slightly lower when accounting for all errors below the current precision limit.



(a) Execution break down for SPHYNX+SBRExecution break down for SPHYNX+SPR by computational workflow step within a single time-step (EC) step (WB)

Figure 6: Break down by workflow step of the time needed to complete one time-step for the EC test (a) and WB test (b) with the resilient code (SPH-YNX+SPR) and the original version (SPHYNX).

To assess the performance of SPR at scale, strong-scaling experiments were conducted on the EC and WB test cases, as described in Table 1.

Each experiment was performed by running two versions of SPHYNX, the original version (SPHYNX) and the resilient version (SPHYNX+SPR). As the focus of this experiment is on assessing the performance of error-detection at scale, given the long execution time requirements for such simulations, the number of time-steps of the simulation was set to 20, while a full SPH simulation would typically require several thousands time-steps. For each execution, we have excluded the first 10 time-steps that correspond to the SPH initialization phase and are not representative of the general performance behavior of the SPH code. Overall, the reported execution time of each test case is an average of a total of 100 time-steps for 10 executions.

Figures 5a and 5b present the average execution time needed to complete a single time-step for both SPHYNX and SPHYNX+SPR, for the EC and WB test cases, respectively, when executed on 4 to 256 computing nodes, each with 12 computing cores. In both cases, it can be observed that SPHYNX+SPR yields a low overhead, never exceeding 5% and 10% for the EC and WB test cases, respectively. This corresponds to the case with the largest number of particles per node.

Figures 6a and 6b illustrate a break down of the execution time of SPHYNX+SPR by the main computational workflow steps, including the time to select particles and the total communication time. It can be seen that the algorithm to select particles represents at most 2.6% of the time needed to complete a single time-step. The total cost of detection never exceeds 0.1% of the total execution time and is not shown in these figures.

Overall, both SPHYNX and SPHYNX+SPR scale very well with increasing node count, the main performance limitations being: (1) the sequential implementation of the Build tree step SPHYNX; and (2) the increasing communication cost, which is negligible when executing on 4 computing nodes, but





Figure 7: Average communication to computation ratio (CCR) per time-step for the EC test (a) and the WB test (b) with the resilient code (SPHYNX+SPR) and the original version (SPHYNX).

Figures 7a and Figure 7b show the average communication to computation ratio (CCR) for the EC and WB test simulations, respectively. For both versions of the code, the CCR ratio steadily increases with the number of nodes. Yet, despite requiring multiple additional inter-process communications, the difference with between SPHYNX and SPHYNX+SPR remains very small. The significant increase in overhead on 256 nodes (also shown in Figure 6a and Figure 6b) is mainly due to having too few particles per node.

3.4 Summary and Discussion

The experimental results show that SPR can practically be incorporated into commonly used astrophysical SPH test simulations with minimal changes to the original code, and with low additional overhead, ranging from 1% to 10% for the EC and WB test cases.

SPH simulations are expected to use more particles in production runs (i.e. up to 10^{12}), which results in greater particle counts per process and greater neighbor count per process (i.e. several hundreds) in some scenarios. Consequently, SPR is expected to deliver improved performance on large-scale SPH simulations: an increase in both the number of particles and number of neighbors per particles will lead to fewer selected particles, and, thus, less overhead overall.

4 Related Work

Complete redundancy The use of redundant MPI processes for error detection has been widely analyzed in the last decade [16, 18, 24, 26, 27, 40, 50]. Unlike these full-process or full-task replication efforts, the present work employs selective *sub-process* or *sub-task* duplication (illustrated in Figure 2) of a selected

part of the work assigned to a process or a task that corresponds to the particle selected for replication. While dual modular redundancy or triple modular redundancy incurs 100% and 200% additional computational overhead, respectively, SPR only replicates 1 - 10% of the application with a 1 - 10% added overhead.

Partial redundancy Partial redundancy has been studied to decrease the overhead of complete redundancy [22, 46–48]. Adaptive partial redundancy has also been proposed wherein a subset of processes is dynamically selected for replication [30]. Partial replication (using additional hardware) of selected MPI processes has been combined with prediction-based detection to achieve SDC protection levels comparable with those of full duplication [9, 10, 36]. The proposed SPR approach differs from the existing partial redundancy approaches in that it protects the data of the entire application (as opposed to a subset) by selectively duplicating *only a subset* of the computations *within* processes. To the best of our knowledge, this is the first time such methods are applied to SPH.

Detection via Interpolation Silent error detectors based on *data analytics* use several interpolation techniques, such as time series prediction [8] and multivariate interpolation [2, 5, 6], to *interpolate* the next value of a computational point based on spatial and temporal data smoothness. Multivariate interpolation has been used to detect and correct SDCs in computational fluid dynamics [5], where the approach achieves 99.8% precision and low recall. Additional tests on synthetic benchmarks achieved a maximum of 90% recall at 1% overhead [6]. To ensure that the false-positive rate remains well below the SDC rate, a high detection precision supersedes a high recall. As shown in Section 3, SPR yields 100% precision (i.e. no false-positives) and a high recall for selected data-fields.

Algorithm-based fault tolerance (ABFT) is another error detection technique which uses checksums to detect up to a certain number of errors and is currently only suitable for linear algebra kernels [11, 32, 42].

5 Conclusion and Future Work

In this work, we proposed a novel silent data corruption (SDC) detection method, namely selective particle replication (SPR), and we implemented it in SPHYNX, a production smoothed particle hydrodynamics (SPH) simulation code.

SPR comprises of three simple, yet efficient, algorithms to select the particles for which to replicate the computations and data of selected particles onto different processes, and to compare the values of the original particles against the data of their replica, to detect SDCs during execution.

The error-detection capabilities of SPR were experimentally evaluated through an error-injection and detection campaign. Not only is SPR able to detect 91 - 99.9% of the injected errors (i.e., DRAM bit-flips), it also does not raise any false-positives. In addition, experiments conducted on an HPC system demonstrate the scalability of SPR in production SPH codes, at an overhead of 1 - 10%. Because SPR is scalable, non-intrusive and precise, it can easily be combined with other error-detection methods, to detect errors that could escape SPR's detection coverage. Furthermore, SPR can also be applied to other classes of applications, e.g., N-body simulations, stencils, and computational fluid dynamics. Whether SPR can accurately detect errors in such applications is a topic that deserves further investigation.

Acknowledgement

This work has been supported by the Swiss Platform for Advanced Scientific Computing (PASC) project SPH-EXA and by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID c16.

References

- [1] O. Agertz, B. Moore, J. Stadel, D. Potter, F. Miniati, J. Read, L. Mayer, A. Gawryszczak, A. Kravtsov, Å. Nordlund, F. Pearce, V. Quilis, D. Rudd, V. Springel, J. Stone, E. Tasker, R. Teyssier, J. Wadsley, and R. Walder. Fundamental differences between SPH and grid methods. *Monthly Notices* of the Royal Astronomical Society, 380:963–978, Sept. 2007.
- [2] L. Bautista Gomez and F. Cappello. Detecting silent data corruption through data dynamic monitoring for scientific applications. In Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), pages 381–382, New York, NY, USA, 2014. ACM.
- [3] L. Bautista-Gomez, S. Tsuboi, D. Komatitsch, F. Cappello, N. Maruyama, and S. Matsuoka. FTI: High performance fault tolerance interface for hybrid systems. In 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC), pages 1–12, Nov 2011.
- [4] L. Bautista-Gomez, F. Zyulkyarov, O. Unsal, and S. McIntosh-Smith. Unprotected computing: A large-scale study of DRAM raw error rate on a supercomputer. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 55:1–55:11, Piscataway, NJ, USA, 2016. IEEE Press.
- [5] L. A. Bautista-Gomez and F. Cappello. Detecting and correcting data corruption in stencil applications through multivariate interpolation. In 2015 IEEE International Conference on Cluster Computing, CLUSTER 2015, Chicago, IL, USA, September 8-11, 2015, pages 595-602, 2015.
- [6] L. A. Bautista-Gomez and F. Cappello. Exploiting spatial smoothness in HPC applications to detect silent data corruption. In 17th IEEE International Conference on High Performance Computing and Communications, HPCC 2015, 7th IEEE International Symposium on Cyberspace Safety and Security, CSS 2015, and 12th IEEE International Conference on Embedded Software and Systems, ICESS 2015, New York, NY, USA, August 24-26, 2015, pages 128–133, 2015.

- [7] A. Benoit, A. Cavelan, Y. Robert, and H. Sun. Optimal resilience patterns to cope with fail-stop and silent errors. In 2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS), pages 202–211, May 2016.
- [8] E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello. Lightweight silent data corruption detection based on runtime data analysis for HPC applications. In *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '15, pages 275–278, New York, NY, USA, 2015. ACM.
- [9] E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello. Exploring partial replication to improve lightweight silent data corruption detection for HPC applications. In Euro-Par 2016: Parallel Processing - 22nd International Conference on Parallel and Distributed Computing, Grenoble, France, August 24-26, 2016, Proceedings, pages 419–430, 2016.
- [10] E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello. Toward general software level silent data corruption detection for parallel applications. *IEEE Trans. Parallel Distrib. Syst.*, 28(12):3642–3655, 2017.
- [11] G. Bosilca, R. Delmas, J. Dongarra, and J. Langou. Algorithm-based fault tolerance applied to high performance computing. *Journal of Parallel and Distributed Computing*, 69(4):410–416, 2009.
- [12] R. M. Cabezón, D. García-Senz, and J. Figueira. SPHYNX: an accurate density-based SPH method for astrophysical applications. Astronomy & Astrophysics, 606:A78, Oct. 2017.
- [13] R. M. Cabezón, D. García-Senz, and A. Relaño. A one-parameter family of interpolating kernels for smoothed particle hydrodynamics studies. *Journal* of Computational Physics, 227:8523–8540, Oct. 2008.
- [14] F. Cappello, A. Geist, B. Gropp, L. Kale, B. Kramer, and M. Snir. Toward exascale resilience. *Int. J. High Perform. Comput. Appl.*, 23(4):374–388, Nov. 2009.
- [15] F. Cappello, A. Geist, W. Gropp, S. Kale, B. Kramer, and M. Snir. Toward exascale resilience: 2014 update. *Supercomputing Frontiers and Innovations*, 1(1), 2014.
- [16] H. Casanova, Y. Robert, F. Vivien, and D. Zaidouni. On the impact of process replication on executions of large-scale parallel applications with coordinated checkpointing. *Future Generation Comp. Syst.*, 51:7–19, 2015.
- [17] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. ACM Transactions on Computer Systems, 3(1):63-75, 1985.
- [18] S. P. Crago, D. I. Kang, M. Kang, R. Kost, K. Singh, J. Suh, and J. P. Walters. Programming models and development software for a space-based many-core processor. In 4th Int. Conf. on Space Mission Challenges for Information Technology, pages 95–102. IEEE, 2011.

- [19] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Comp. Syst.*, 22(3):303–312, 2006.
- [20] W. Dehnen and H. Aly. Improving convergence in smoothed particle hydrodynamics simulations without pairing instability. *Monthly Notices of* the Royal Astronomical Society, 425:1068–1082, Sept. 2012.
- [21] S. Di, Y. Robert, F. Vivien, and F. Cappello. Toward an optimal online checkpoint solution under a two-level HPC checkpoint model. *IEEE Transactions on Parallel and Distributed Systems*, 28(1):244–259, Jan 2017.
- [22] J. Elliott, K. Kharbas, D. Fiala, F. Mueller, K. Ferreira, and C. Engelmann. Combining partial redundancy and checkpointing for HPC. In 2012 IEEE 32nd International Conference on Distributed Computing Systems (ICDCS), pages 615–626, June 2012.
- [23] E. N. M. Elnozahy, L. Alvisi, Y.-M. Wang, and D. B. Johnson. A survey of rollback-recovery protocols in message-passing systems. ACM Computing Survey, 34:375–408, 2002.
- [24] C. Engelmann and S. Boehm. Redundant execution of HPC applications with MR-MPI. In Proceedings of the 10th IASTED International Conference on Parallel and Distributed Computing and Networks, PDCN 2011, 04 2011.
- [25] A. E. Evrard. Beyond N-body 3D cosmological gas dynamics. Monthly Notices of the Royal Astronomical Society, 235:911–934, Dec. 1988.
- [26] K. Ferreira, J. Stearley, J. H. Laros, R. Oldfield, K. Pedretti, R. Brightwell, R. Riesen, P. G. Bridges, and D. Arnold. Evaluating the viability of process replication reliability for exascale systems. In *Conference for High Perfor*mance Computing, Networking, Storage and Analysis (SC), pages 1–12, Nov 2011.
- [27] D. Fiala, F. Mueller, C. Engelmann, R. Riesen, K. B. Ferreira, and R. Brightwell. Detection and correction of silent data corruption for largescale high-performance computing. In *Conference on High Performance Computing Networking, Storage and Analysis (SC)*, page 78, 2012.
- [28] D. García-Senz, R. M. Cabezón, and J. A. Escartín. Improving smoothed particle hydrodynamics with an integral approach to calculating gradients. *Astronomy & Astrophysics*, 538:A9, Feb. 2012.
- [29] M. R. Garey and D. S. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York, NY, USA, 1990.
- [30] C. George and S. S. Vadhiyar. ADFT: An adaptive framework for fault tolerance on large scale systems using application malleability. *Proceedia Computer Science*, 9:166 – 175, 2012.
- [31] R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics - Theory and application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society*, 181:375–389, Nov. 1977.

- [32] K.-H. Huang and J. A. Abraham. Algorithm-based fault tolerance for matrix operations. *IEEE Trans. Comput.*, 33(6):518–528, 1984.
- [33] A. A. Hwang, I. A. Stefanovici, and B. Schroeder. Cosmic Rays Don't Strike Twice: Understanding the Nature of DRAM Errors and the Implications for System Design. In Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XVII, pages 111–122, New York, NY, USA, 2012. ACM.
- [34] L. B. Lucy. A numerical approach to the testing of the fission hypothesis. Astronomical Journal, 82:1013–1024, Dec. 1977.
- [35] R. E. Lyons and W. Vanderkulk. The use of triple-modular redundancy to improve computer reliability. *IBM J. Res. Dev.*, 6(2):200–209, 1962.
- [36] X. Ni and L. V. Kale. FlipBack: Automatic Targeted Protection against Silent Data Corruption. In Conference for High Performance Computing, Networking, Storage and Analysis (SC), pages 335–346, Nov 2016.
- [37] B. Nie, D. Tiwari, S. Gupta, E. Smirni, and J. H. Rogers. A large-scale study of soft-errors on GPUs in the field. In 2016 IEEE International Symposium on High Performance Computer Architecture, pages 519–530, March 2016.
- [38] G. Oger, D. L. Touzé, D. Guibert, M. de Leffe, J. Biddiscombe, J. Soumagne, and J.-G. Piccinali. On distributed memory MPI-based parallelization of sph codes in massive HPC context. *Computer Physics Communications*, 200:1 – 14, 2016.
- [39] T. O'Gorman. The effect of cosmic rays on the soft error rate of a DRAM at ground level. *IEEE Trans. Electron Devices*, 41(4):553–557, 1994.
- [40] M. W. Rashid and M. C. Huang. Supporting highly-decoupled thread-level redundancy for parallel programs. In Conf. on High-Performance Computer Architecture (HPCA), pages 393–404. IEEE, 2008.
- [41] S. Rosswog. Boosting the accuracy of SPH techniques: Newtonian and special-relativistic tests. *Monthly Notices of the Royal Astronomical Soci*ety, 448:3628–3664, Apr. 2015.
- [42] M. Shantharam, S. Srinivasmurthy, and P. Raghavan. Fault tolerant preconditioned conjugate gradient for sparse linear system solution. In *ICS*. ACM, 2012.
- [43] M. Snir, R. W. Wisniewski, J. A. Abraham, S. V. Adve, S. Bagchi, P. Balaji, J. Belak, P. Bose, F. Cappello, B. Carlson, A. A. Chien, P. Coteus, N. A. Debardeleben, P. C. Diniz, C. Engelmann, M. Erez, S. Fazzari, A. Geist, R. Gupta, F. Johnson, S. Krishnamoorthy, S. Leyffer, D. Liberty, S. Mitra, T. Munson, R. Schreiber, J. Stearley, and E. V. Hensbergen. Addressing failures in exascale computing. *Int. J. High Perform. Comput. Appl.*, 28(2):129–173, 2014.
- [44] V. Springel. Smoothed Particle Hydrodynamics in Astrophysics. Annual Reviews of Astronomy and Astrophysics, 48:391–430, Sept. 2010.

- [45] V. Sridharan, N. DeBardeleben, S. Blanchard, K. B. Ferreira, J. Stearley, J. Shalf, and S. Gurumurthi. Memory errors in modern systems: The good, the bad, and the ugly. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '15, pages 297–310, New York, NY, USA, 2015. ACM.
- [46] J. Stearley, K. Ferreira, D. Robinson, J. Laros, K. Pedretti, D. Arnold, P. Bridges, and R. Riesen. Does partial replication pay off? In *IEEE/IFIP International Conference on Dependable Systems and Networks Workshops* (DSN 2012), pages 1–6, June 2012.
- [47] O. Subasi, J. Arias, O. Unsal, J. Labarta, and A. Cristal. Programmerdirected partial redundancy for resilient hpc. In *Proceedings of the 12th* ACM International Conference on Computing Frontiers (CF), pages 47:1– 47:2, New York, NY, USA, 2015. ACM.
- [48] O. Subasi, G. Yalcin, F. Zyulkyarov, O. Unsal, and J. Labarta. Designing and Modelling Selective Replication for Fault-Tolerant HPC Applications. In 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), pages 452–457, May 2017.
- [49] J. W. Young. A first order approximation to the optimum checkpoint interval. Comm. of the ACM, 17(9):530–531, 1974.
- [50] J. Yu, D. Jian, Z. Wu, and H. Liu. Thread-level redundancy fault tolerant cmp based on relaxed input replication. In 2011 6th International Conference on Computer Sciences and Convergence Information Technology (ICCIT), pages 544–549, Nov 2011.
- [51] J. Ziegler, M. Nelson, J. Shell, R. Peterson, C. Gelderloos, H. Muhlfeld, and C. Montrose. Cosmic ray soft error rates of 16-Mb DRAM memory chips. *IEEE Journal of Solid-State Circuits*, 33(2):246–252, 1998.