

Proactive Process-Level Live Migration in HPC Environments

Chao Wang¹, Frank Mueller¹, Christian Engelmann², Stephen L. Scott²

¹ Department of Computer Science, North Carolina State University, Raleigh, NC

² Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN
mueller@cs.ncsu.edu, phone: +1.919.515.7889, fax: +1.919.515.7896

Abstract

As the number of nodes in high-performance computing environments keeps increasing, faults are becoming common place. Reactive fault tolerance (FT) often does not scale due to massive I/O requirements and relies on manual job resubmission.

This work complements reactive with proactive FT at the process level. Through health monitoring, a subset of node failures can be anticipated when one's health deteriorates. A novel process-level live migration mechanism supports continued execution of applications during much of processes migration. This scheme is integrated into an MPI execution environment to transparently sustain health-inflicted node failures, which eradicates the need to restart and requeue MPI jobs. Experiments indicate that 1-6.5 seconds of prior warning are required to successfully trigger live process migration while similar operating system virtualization mechanisms require 13-24 seconds. This self-healing approach complements reactive FT by nearly cutting the number of checkpoints in half when 70% of the faults are handled proactively.

I. Introduction

The current trend in high-performance computing (HPC) is to continually increase the system scale through ever larger number of nodes, each of them consisting of multi-core processors, exceeding the 100,000 processor mark. This substantial system growth poses a challenge in terms of the failure probability for large-scale jobs.

Reliability data of contemporary systems, depicted in Table I, illustrates that the mean time between failures (MTBF) / interrupts (MTBI) is in the range of 6.5-40 hours, depending on the maturity / age of the installation [22]. The most common causes of failure are processor, memory and storage errors / failures. When extrapolating for current systems, the mean-time between failures (MTBF) for peta-scale systems is predicted to be as short as 1.25 hours [35] and could be accompanied by one radiation-induced soft error once every 1.5 hours [38]. With current checkpoint/restart (C/R) technology, the wall-clock time of a 100-hour job could well extend to 251

hours due to C/R overhead implying that 60% of cycles are spent on C/R alone [35].

System	# CPUs	MTBF/I
ASCI Q	8,192	6.5 hrs
ASCI White	8,192	5/40 hrs ('01/'03)
PSC Lemieux	3,016	9.7 hrs
Google	15,000	20 reboots/day

TABLE I: Reliability of HPC Clusters

Yet, when fault-tolerant middleware hides such failures and user services remain unaffected, the interpolated fault rate is just over one hour for equivalent number of nodes (c.f. Google in Table I) [18]. In this spirit, our work focuses on fault-tolerant middleware for HPC systems. More specifically, this paper promotes process-level live migration combined with health monitoring for a proactive FT approach that complements existing C/R schemes with self healing whose fault model is based on the work by Tikotekar *et al.* [48].

Health monitoring has recently become a wide-spread feature in commodity and, even more so, in server and HPC components. Such monitors range from simple processor temperature sensors to baseboard management cards (BMCs) with a variety of sensing capabilities, including fan speeds, voltage levels and chassis temperatures. Similarly, the SMART disk standard provides the means to trigger events upon disk errors indicating disk problems, which can be saved in log files or which can trigger exception handling mechanisms. Aided by such capabilities, node failures may be anticipated when the health status of a node deteriorates, *i.e.*, when abnormal sensor readings or log entries are observed.

Health monitoring has been used to model failure rates and, in a reactive manner, to determine checkpoint intervals [42], [39]. In this work, we venture beyond reactive schemes by promoting a *proactive* approach that migrates processes away from “unhealthy” nodes to healthy ones. Such a self-healing approach has the advantage that checkpoint frequencies can be reduced as sudden, unexpected faults should become the exception. This requires the availability of spare nodes, which is becoming common place in recent cluster acquisitions. We expect such spare nodes to become a commodity provided by job schedulers upon request. Our experiments assume availability of 1-2

This work was supported in part by NSF grants CCR-0237570 (CA-REER), CNS-0410203, CCF-0429653 and DOE DE-FG02-05ER25664. The research at ORNL was supported by Office of Advanced Scientific Computing Research and DOE DE-AC05-00OR22725 with UT-Battelle, LLC.

spare nodes.¹

The feasibility of health monitoring at various levels has recently been demonstrated for temperature-aware monitoring, *e.g.*, by using ACPI [1], and, more generically, by critical-event prediction [40]. Particularly in systems with thousands of processors, fault handling becomes imperative, yet approaches range from application-level and runtime-level to the level of OS schedulers [8], [7], [9], [34]. These and other approaches differ from our work in that we promote *live migration combined with health monitoring*.

We have designed an automatic and transparent mechanism for *proactive* FT of arbitrary MPI applications. The implementation, while realized over LAM (Local Area Multicomputer)/MPI’s C/R support [43] through Berkeley Labs C/R (BLCR) [15], is in its mechanisms applicable to any process-migration solution, *e.g.*, the Open MPI FT mechanisms [24], [25]. BLCR is an open source, system-level C/R implementation integrated with LAM/MPI via a callback function. The original LAM/MPI+BLCR combination [41] only provides reactive FT and requires a complete job restart from the last checkpoint including job resubmission in case of a node failure. Recent work enhances this capability with a job pause/continue mechanism that keeps an MPI job alive while a failed node is replaced by a spare node [51]. Paused, healthy tasks are rolled back to and spare nodes proceed from the last checkpoint in a coordinated manner transparent to the application.

The contribution of this paper is to avoid roll-backs to prior checkpoints whenever possible. By monitoring the health of each node, a process is migrated as a precaution to potentially imminent failure. To reduce the impact of migration on application performance, we contribute a novel *process-level live migration* mechanism as an enhancement to the Linux BLCR module. Thus, execution proceeds while a process image is incrementally and asynchronously transferred to a spare node. This reduces the time during which the process is unresponsive to only a short freeze phase when final changes are transferred to the spare node before re-activating execution on the target node. Hence, MPI applications execute during much of process migration. In experiments, we assessed the trade-off between lower end-to-end wall-clock times of jobs subject to live migration *vs.* the slightly prolonged duration for migration as opposed to a traditional process-freeze approach. Depending on the estimated remaining up-time of a node with deteriorating health, one can choose

¹Our techniques also generalize to task sharing on a node should not enough spare nodes be available, yet the cost is reduced performance for tasks on such shared nodes. This may result in imbalance between all tasks system-wide resulting in decreased overall performance. Such imbalance might be tolerable when faulty nodes can be brought back online quickly so that processes can migrate back to their original nodes.

between live and frozen migration schemes.

Our results further demonstrate that proactive FT complements reactive schemes for long-running MPI jobs. Specifically, should a node fail without prior health indication or while proactive migration is in progress, our scheme reverts to reactive FT by restarting from the last checkpoint. Yet, as proactive FT has the potential to prolong the mean-time-to-failure, reactive schemes can lower their checkpoint frequency in response, which implies that proactive FT can lower the cost of reactive FT. More specifically, experimental results indicate that 1-6.5 seconds of prior warning are required to successfully trigger live process migration while similar operating system (OS) virtualization mechanisms require 13-24 seconds. The approach further complements reactive FT by nearly twice as long a checkpointing interval due to proactive migration when 70% of the failures are predicted only a few seconds prior (derived from [40]).

II. Design

Figure 1 depicts the system components and their interaction, *i.e.*, the chronological steps involved in process migration of each MPI job and job dependencies with data exchanges. In the following, we discuss system support for live migration at two levels: (1) the synchronization and coordination mechanisms within an MPI job and (2) live migration with incremental update support at the process/MPI task level. We further consider the tradeoff between live and frozen migration options and develop approaches to manage “dirty” memory, *i.e.*, memory written since the last incremental update.

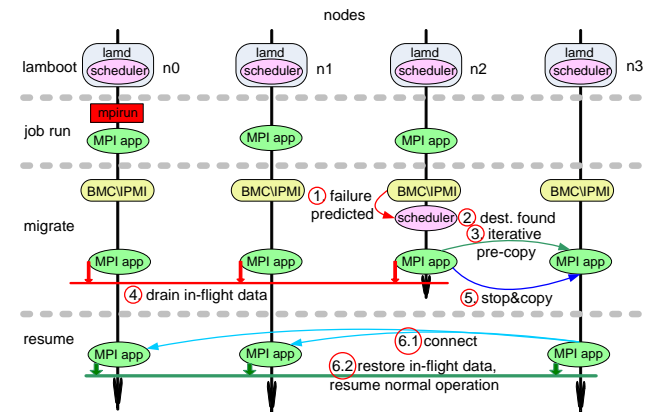


Fig. 1: Job Live Migration

A. Live Migration at the Job Level

Figure 1 depicts steps 1-6, each of which are described in the following.

Step 1: Migration Trigger: In our system, the per-node health monitoring mechanism is realized on top of a Baseboard Management Controller (BMC). It is equipped with sensors to monitor different properties, *e.g.*, sensors providing data on temperature, fan speed, and voltage. We

also employ the Intelligent Platform Management Interface (IPMI), an increasingly common management/monitoring component that provides a standardized message-based interface to obtain sensors readings for health monitoring. We further designed a decentralized scheduler, which can be deployed as a stand-alone component or as an integral process of an MPI daemon, such as the LAM daemon (lamd). The scheduler will be notified upon deteriorating health detected by BMC/IPMI, *e.g.*, due to a sensor reading exceeding a threshold value.

Step 2: Destination Node Determination: When the scheduler component on the health-decaying node receives the early warning issued in step 1, it first chooses a spare node as a replacement for the unhealthy node. Spare nodes are increasingly becoming a commodity in clusters. We expect that job schedulers may soon simply allocate a limited number of spares as a default or upon request during job submission. As of now, we still have to explicitly over-provision by setting a constant number of nodes aside during job submission.

These spare nodes comprise the destination for process migration. However, if no spare node was set aside during job submission or if all nodes were already used, we choose the most lightly loaded node as a migration destination, in which case the node doubles for two MPI tasks. Sharing nodes may result in imbalance due to bearing twice the load of regular nodes, which generally results in lower overall application performance. Such imbalance might be tolerable when faulty nodes can be brought back online quickly so that processes can migrate back to their original nodes. Nonetheless, higher sustained application performance can be guaranteed when unloaded spares are available as migration targets.

Step 3: Memory Precopy: Once a destination node is determined, the scheduler initiates the migration of a process on both destination and source nodes. The objective of the first stage of migration is to transfer a memory snapshot of the process image to the spare node, yet to allow the application to execute during this stage, hence the name *live* migration. The transfer of the process image occurs at page granularity and is repeated for pages written to by the application between image scans. During the first process image scan (first iteration), all non-zero pages are transferred from the source node to the destination node. On subsequent scans/iterations, only the pages updated since the previous scan are transferred. When the number of such “dirty” pages between scans does not change significantly anymore, the scan loop terminates. System support for tracking dirty pages is discussed later in the context of memory management.

Step 4: In-flight Message Drainage: Before we stop the process and migrate the remaining dirty pages with the corresponding process state to the destination node, all

MPI tasks need to coordinate to reach a consistent global state. Based on our LAM/MPI+BLCR design, message passing is dealt with at the MPI level while the process-level BLCR mechanism is not aware of messaging at all. Hence, we employ LAM/MPI’s job-centric interaction mechanism for the respective MPI tasks to clear in-flight data in the MPI communication channels.

Step 5: Stop&Copy: Once all the MPI tasks (processes) reach a consistent global state, the process on the source node freezes (suspends) application execution but still copies the remaining dirty pages (written to since the last iteration in step 3) and the final process state (registers, signal information, pid, files etc.) to the destination node. All other MPI tasks are suspended at their point of execution.

Step 6: Connection Recreation, Messages Restoration and Job Continuation: When the process is ready on the destination node, it sets up a communication channel with all other MPI tasks. Subsequently, the drained in-flight messages are restored, and all the processes resume execution from the point of suspension.

B. Live Migration at the Process Level

The incremental precopy and stop© (steps 3 and 5 in Figure 1) are performed at the process level involving only the destination and source nodes. Yet, there are trade-offs between simply stopping an application to engage in a frozen copy and the alternative of a live precopy with continued application progress. The latter, while generally resulting in shorter overall application wall-clock time, comes at the expense of background network activity, possibly even repeatedly transmitting dirtied pages. This also raises the question when the iterative precopy loop should terminate.

Figures 2 and 3 show migration with precopy (live) and without (stop©-only). Compared to a traditional job resubmission resuming execution from the last checkpoint in a reactive FT scheme, both of these proactive migration schemes lower the expected execution time of the MPI application in the presence of failures. This is especially the case for HPC environments where the MTBF is low, which typically implies that the number of compute nodes is high, the run-time of an application is long, and the memory footprint is large.

One of our objectives is to reduce the aggregate downtimes over all nodes, *i.e.*, the duration of the stop© step should be small. Live migration with incremental precopy results not only in shorter downtime on the local node (to transfer dirty pages plus other process state) but also in reduced suspension of all other nodes (once MPI message queues are drained) since fewer pages remain dirty after precopy. Another objective is to tolerate the imminent fault. The shorter the migration duration, the higher the probability that our proactive scheme makes

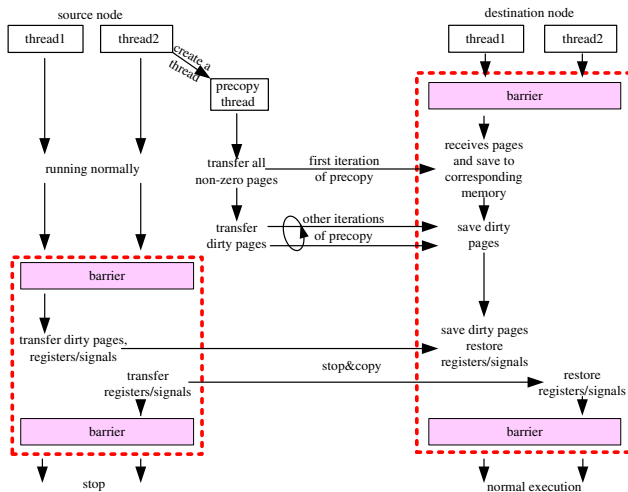


Fig. 2: Process Migration with Precopy (Kernel Mode in Dotted Frame)

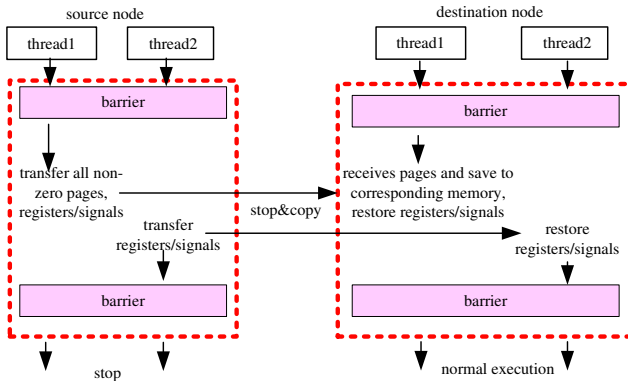


Fig. 3: Process Migration without Precopy (Kernel Mode in Dotted Frame)

a reactive restart from a prior checkpoint unnecessary. Frozen migration consisting only of the stop© step takes less overall time during migration, thereby increasing chances for successful migration. A compromise might even be to stop the precopy step prematurely upon receipt of another fault event indicating higher urgency of the health problem.

Two major factors affect the tradeoff between the downtime and the migration duration. First, the network bandwidth shared between the MPI job and the migration activity is limited. If an application utilizes less network bandwidth, more bandwidth can be consumed by the migration operation. Thus, the precopy step may only have minimum impact on application progress. More communication-intensive applications, however, may leave less unused bandwidth for live migration (during precopy). This could both slow down a local MPI task (with potentially globally unbalanced progress due to communication) and prolong precopy duration, ultimately requiring a premature termination of this step. High-end HPC installations may provide multiple network backbones to isolate MPI communication from I/O traffic, the latter of which

covering checkpointing and migration traffic as well. In such systems, bandwidth issues may not be as critical (but cannot be fully discounted either).

Second, the page write rate (dirtying rate) affects the trade-off between downtime and migration duration. The dirtying rate is affected by the memory footprint (more specifically, the *rewrite* working set) of the application. A larger number of pages repeatedly written within a tight loop will prolong the precopy duration, which might lower the probability for successful migration (compared to non-live migration). In fact, the page access patterns of the studied benchmarks will be shown to differ significantly. We further discuss this issue in the experimental section.

C. Memory Management

A prerequisite of live migration is the availability of a mechanism to track modified pages during each iteration of the precopy loop. Two fundamentally different approaches may be employed, namely page protection mechanisms or page-table dirty bits. Different implementation variants build on these schemes, such as the bookkeeping and saving scheme that, based on the dirty bit scheme, copies pages into a buffer [19].

Under the page protection scheme, all writable pages of a process are write-protected before each iteration occurs. When a page is modified (written to for the first time), a page fault exception is raised. The triggered exception handler enables write access to the page and records the page as dirty. Upon return from the handler, the write access is repeated successfully, and subsequent writes to the same page execute at full speed (without raising an exception). At the next iteration, pages marked as dirty are (1) write protected again before being (2) copied to the destination node. This ordering of steps is essential to avoid races between writes and page transfers.

Under the dirty bit scheme, the dirty bit of the page-table entry (PTE) is duplicated. The dirty bit is set in response to handling by the memory management unit (MMU) whenever the first write to the page is encountered. At each iteration, the duplicate bit is checked and, if set, is first cleared before the page is transferred to the destination node. To provide this shadow functionality, kernel-level functions accessing the PTE dirty bit are extended to also set the duplicate bit upon the first write access.

The page protection scheme has certain draw-backs. Some address ranges, such as the stack, can only be write protected if an alternate signal stack is employed, which adds calling overhead and increases cache pressure. Furthermore, the overhead of user-level exception handlers is much higher than kernel-level dirty-bit shadowing. Thus, we selected the dirty bit scheme in our design. Implementation-specific issues of this scheme are discussed in detail in the next section.

III. Implementation

We next provide implementation details of our process-level live migration under LAM/MPI+BLCR. The overall design and the principle implementation methodology are applicable to arbitrary MPI implementations, *e.g.*, OpenMPI and MPICH.

A. Failure Prediction and Decentralized Scheduler

As outlined in Section II, the capabilities of the BMC hardware and IPMI software abstraction were utilized by our framework to predict failures. In the following, we focus on the system aspects of the live migration mechanism to tolerate imminent faults and only to a lesser extent on fault prediction, a topic mostly beyond the scope of this paper.

The scheduler is implemented as a process of `lamd` communicating through the out-of-band channel provided by LAM/MPI. When the scheduler daemon is notified by the BMC/IPMI health monitoring component, it consults the database to retrieve information about the MPI jobs and the nodes in the LAM universe. A spare node or, alternatively, the most lightly loaded node is chosen as the migration target. In our implementation, we also check the BMC/IPMI sensor readings with preset thresholds to determine if the degree of urgency allows accommodation of a precopy step. For example, if the temperature is higher than a high watermark (indicating that failure is almost instantly imminent), the frozen migration (`stop©-only`) is launched, which guarantees a shorter migration duration. However, if the value is closer to the low watermark (indicating that failure is likely approaching but some time remains), the live migration (with its precopy step) is launched, which reduces overall application downtime. The design allows additional schemes, selected through a combination of sensor and history information, to be integrated as plug-ins in the future.

Next, we present implementation details of the process-level mechanisms, including dirty page tracking and process image restoration, as they were realized within BLCR. This is followed by the MPI-level implementation details based on the fundamental capabilities of BLCR, including the maintenance of consistency of a global state for the entire MPI job, as realized within LAM/MPI.

B. Process-Level Support for Live Migration

As part of our implementation, we integrated several new BLCR features to extend its process-level task migration facilities and to coordinate LAM/MPI's callback mechanism amongst all MPI processes. The scheduler discussed previously issues the `cr_save` and `cr_restore` commands on the source and destination nodes, respectively, when it deems live migration with precopy to be beneficial. Subsequently, once the precopy step has completed, LAM/MPI issues the `cr_stop` and `cr_quiesce` commands on the source node and all other operational nodes, respectively.

1) *Precopy at the Source Node: cr_save*: The `cr_save` command implements a sequence of steps specific to our live migration. It first sends a signal to the process on the source node where it triggers the creation of a new thread. This thread performs the precopy work at the user level. We chose a solution at the user level since a kernel-level precopy can block application threads, thereby hampering overall computational progress. During the first iteration of the precopy loop, all non-empty pages are transferred to the destination node. Subsequent iterations, in contrast, result in transfers of only those pages modified since the last iteration. The top half of Figure 2 at the source node depicts this procedure. Recall that the dirty bit approach implemented in our system tracks if a page has been written to since the last transfer (iteration).

The Linux kernel maintains a dirty bit in the PTE. It is automatically set by the hardware (MMU) when the page is modified. Yet, we cannot utilize this dirty bit for our purposes since its state is tightly coupled with the Linux memory management, specifically the swap subsystem. We have to clear such a flag for each iteration of our page sweep, yet clearing the PTE dirty bit would indicate that a page need not be written on a subsequent swap-out, which is violating consistency. Instead of modifying the dirty bit semantics, we decided to shadow the dirty bit within the reserved bits of PTE. Updates to this shadow dirty bit may occur through system functions without much overhead (compared to user-level page protection).

This shadow dirty bit still needs to be preserved across swapping, which is a challenge since swapped PTEs are partially invalidated. One solution would be to simply disable swapping, which might be an option for HPC, but a more general solution is preferable. Linux actually preserves selected fields of a swapped out PTE, among other things to associate with it the swap disk device and the offset into the device where the page resides. This swap information includes a set of bits that are preserved, one of which we utilized as the shadow dirty bit. We implemented this approach in `x86_64` and `i386` Linux kernels similar to an orthogonal IA64 design by HP [2].

During live migration, a precopy thread iteratively transfers dirty pages to the destination node until one of the following conditions is satisfied:

- The aggregate size of dirty memory during the last iteration is less than a lower *memory threshold* (default: 1MB);
- the aggregate difference of transferred memory in consecutive iterations is less than a small *difference threshold* (indicating that a steady state of the rate in which pages are written to is reached);
- the remaining time left before the expected failure is below a lower *overhead threshold*.

As shown in Figures 4(i) and 4(j), the page dirty mod-

ification eventually stabilizes or its fluctuation is regular (repeats periodically). We could support an empirical factor in the algorithm to keep a profile history of the page modification rate and its regularity. Such model parameters could steer future application runs during live migration / precopy in choosing a more sensitive termination condition. This could be especially beneficial for jobs with long runtime (a large number of timesteps) or repeated job invocations.

Once the precopy step terminates, the thread informs the scheduler, who then enters the stop© step in a coordinated manner across all processes of the MPI job. This includes issuing the *cr_stop* command on the source node.

2) *Freeze at the Source Node: cr_stop*: The redesigned *cr_stop* command signals the process on the source node to freeze execution, *i.e.*, to stop and copy the pages dirtied in the last iteration of the precopy step to the destination node. Threads of the process subsequently take turns copying their own state information (registers, signal information, etc.) to the destination node. This functionality is performed inside the Linux kernel (as an extension to BLCR) since process-internal information needs to be accessed directly and in a coordinated fashion between threads. The lower half of Figure 2 depicts these actions at the source node.

Between the precopy and freeze steps, the processes of an MPI job also need to be globally coordinated. The objective is to reach a consistent global state by draining all in-flight messages and delaying any new messages. This is accomplished in two parts at the source node, implemented by the *cr_save* and *cr_stop* commands. At the destination node, a single command implements the equivalent restore functionality.

3) *Precopy and Freeze at the Destination Node: cr_restore*: At the destination node, the *cr_restore* command with our extensions is issued by the scheduler. This results in the immediate creation of an equal number of threads as were existent on the source node, which then wait inside the Linux kernel for messages from the source node. A parameter to the command, issued by the scheduler, indicates whether or not a live/precopied step was selected on the source node. In either case, one thread receives pages from the source node and places them at the corresponding location in local memory. All threads subsequently restore their own state information received from the corresponding source node threads, as depicted in Figures 2 and 3 at the destination node. After the process image is fully transferred and restored, the user mode is entered. Next, the MPI-level callback function is triggered, which creates the connections with the other MPI processes of the compute job and restores the drained in-flight messages discussed next.

4) *Process Quiesce at Operational Nodes: cr_quiesce*: In steps 4 and 6 of our design (cf. Section II), processes on all other operational nodes drain the in-flight messages before the stop© step. They then remain suspended in this step, creating a connection with the new process on the destination (spare) node and ultimately restoring in-flight messages after. This sequence of actions is triggered through the newly developed *cr_quiesce* command. In our implementation, issuing this command signals the process (MPI task), which subsequently enters the MPI-level callback function to drain the messages, waits for the end of the stop© step of the faulty/spare nodes, and then restores its communication state before resuming normal execution.

C. Job Communication and Coordination Mechanism for Live Migration

In our implementation, we extended LAM/MPI with fundamentally new functionality provided through our BLCR extension to support live migration of a process (MPI task) within a job. Our approach can be divided into four stages introduced in their respective temporal order:

Stage 1: Live/Precopy: The scheduler of the LAM daemon, *lamd*, determines whether or not live/precopied should be triggered. If sufficient time exists to accommodate this stage, the *cr_save* and *cr_restore* commands are issued on the source and destination nodes, respectively. During the precopy step, the compute job continues to execute while dirty pages are sent iteratively by a precopy thread on the source node. If time does not suffice to engage in this live/precopied stage, the next stage is entered immediately.

Stage 2: Pre-Stop&Copy: In this stage, the scheduler issues a stop© command for *mpirun*, the initial LAM/MPI process at job invocation, which subsequently invokes the *cr_stop* and *cr_quiesce* commands on the source node and any other operational nodes, respectively. Once signaled, any of these processes first enters the callback function, which had been registered as part of the LAM/MPI initialization at job start. The callback forces a drain of the in-flight data to eventually meet a consistent global state for all processes (MPI tasks).

Stage 3: Stop&Copy: Upon return from the callback function, the process on the source node stops executing and transfers the remaining dirty pages and its process state. Meanwhile, the processes on other operational nodes suspend in a waiting pattern, as discussed previously.

Stage 4: Post-Stop&Copy: Once all the state is transferred and restored at the destination (spare) node, the process is activated at this node and invokes the LAM/MPI callback function again from the signal handler. Within the callback handler, the migrated process relays its addressing information to *mpirun*, which broadcasts this information to all other processes. These other processes update their entry in the local process list before establishing a con-

nection with the migrated process on the spare. Finally, all processes restore the drained in-flight data and resume execution from the stopped state.

IV. Experimental Framework

Experiments were conducted on a dedicated Linux cluster comprised of 17 compute nodes, each equipped with two AMD Opteron-265 processors (each dual core) and 2 GB of memory. The nodes are interconnected by two networks, both 1 Gbps Ethernet. The OS used is Fedora Core 5 Linux x86_64 with our dirty bit patch as described in Sections II and III. One of the two networks was reserved for LAM/MPI application communication while the other supported process migration and other network traffic. In the next section, we present results for these two network configurations. As discussed in Section II-B, the MPI job and the migration activity may compete for network bandwidth if only a single network is available. However, it is common for high-end HPC to install two separate networks, one reserved for MPI communication and the other for operations such as I/O, booting, system setup and migration. In our experiments, we also assessed the system performance with a single network responsible for both MPI communication and migration activities. The results are not significantly different from those with two networks, which shows that for the applications evaluated in our system communication and memory intensity do not coincide. For the future work, we will create and assess applications with varying communication rates and memory pressure to measure the tradeoff between live and frozen migrations and to provide heuristics accordingly, as discussed in the next section.

We have conducted experiments with several MPI benchmarks. Health deterioration on a node is simulated by notifying the scheduler daemon, which immediately initiates the migration process. To assess the performance of our system, we measure the wall-clock time for a benchmark with live migration, with stop©-only migration and without migration. The migration overheads are introduced by transferring the state of the process, including the dirty pages, and the coordination among the MPI tasks. In addition, the actual live migration duration can be attributed to two parts: (1) the overhead incurred by the iterative precopy and (2) the actual downtime for which the process on the source node is stopped. Accordingly, precopy durations and downtimes are measured.

V. Experimental Results

Experiments were conducted to assess (a) overheads associated with the migration, (b) frozen and live migration durations, (c) scalability for task and problem scaling of migration approaches, and (d) page access patterns and migration times.

Results were obtained for the NAS parallel benchmarks (NPB) version 3.2.1 [52], a suite of programs widely used

to evaluate the performance of parallel systems. Out of the NPB suite, the BT, CG, FT, LU and SP benchmarks were exposed to class B and class C data inputs running on 4, 8 or 9 and 16 nodes. Some NAS benchmarks have 2D, others have 3D layouts for 2^3 or 3^2 nodes, respectively. In addition to the 16 nodes, one spare node is used, which is omitted (implicitly included) in later node counts in this paper. The NAS benchmarks EP, IS and MG were not suitable for our experiments since they execute for too short a period to properly gauge the effect of imminent node failures. (With class B data inputs, completion times for EP, IS and MG are 11.7, 4.8 and 4.1 seconds, respectively, for 16-node configurations.)

A. Migration Overhead

The first set of experiments assesses the overhead incurred due to one migration (equivalent to one imminent node failure). Figure 4(a) depicts the job execution time without any migration, with live migration and with frozen (stop©-only) migration under class C inputs on 16 nodes. The corresponding overheads per scheme are depicted in Figure 4(b). The results indicate that the wall-clock time for execution with live migration exceeds that of the base run by 0.08-2.98% depending on the application. The overhead of frozen migration is slightly higher at 0.09-6%. The largest additional cost of 6.7% was observed for FT under class B inputs for 16 nodes (not depicted here) due to its comparatively large memory footprint (113 MB) and relatively short execution time (37 seconds). Amongst BT, CG, FT, LU and SP under classes B and C running on 4, 8 or 9 and 16 nodes, the longest execution time is 20 minutes (for SP under class C inputs on 4 nodes, also not depicted here). Projecting these results to even longer-running applications, the overhead of migration becomes less significant (if not even insignificant) considering the equivalent checkpointing overhead under current mean-time-to-failure (MTTF) rates.

B. Migration Duration

Besides overhead due to migration, we assessed the actual migration duration including live migration (precopy duration) and the downtime of other nodes during migration schemes. Figure 4(c) depicts (a) live migration with the precopy and downtime fractions and (b) frozen migration with only its downtime during the stop© step, both for NPB with class C inputs on 16 nodes. The precopy duration was measured from the issuing of the *cr_save* command to its completion. The stop© downtime was measured from issuing *cr_stop* on the source node / *cr_quiesce* on the operational nodes to the resumption of the job execution. Live downtime ranged between 0.29-3.87 seconds while stop© downtime was between 1.04-7.84 seconds. Live migration pays a penalty for the shorter downtime in that its precopy duration is prolonged. Precopy adds another 2.35-24.4 seconds. Nonetheless,

the precopy stage does not significantly affect execution progress of the application as it proceeds asynchronously in the background. Figure 4(d) illustrates this fact for each scheme by comparing the downtime (from Figure 4(b)) with migration overhead for both frozen (stop©) and live migration. Both schemes show a close match between their respective downtime and overhead numbers. (Some benchmarks show shorter overhead than the absolute downtime due to larger variances in job execution times subsuming the shorter overheads, see Section V-C.)

Figure 4(e) depicts the amount of memory transferred during migration. With frozen (stop©-only) migration, memory pages of a process cannot be modified while the process remains inactive in this stage. In contrast, live migration allows pages to be modified and consequently requires repeated transfers of dirty pages. Hence, both precopy duration and downtime are a function of the write frequency to disjoint memory pages. Frozen (stop©-only) migration results in larger memory transfers (50.7-448.6MB) than just the stop© step of live migration (1.7-251MB), yet the latter incurs additional transfers (127.4-1565MB) during the precopy step. This result is consistent with the downtime observations of the two schemes discussed above. Our experiments also indicate an approximate cost of 0.3 seconds for MPI-level coordination and communication during live migration plus a cost of less than 0.1 seconds transferring process state information, *e.g.*, registers and signal information. Hence, both the precopy duration and the downtime are almost entirely due to the overhead of transferring memory pages. Hence, the overall trends and patterns of Figures 4(c) and 4(e) tend to be similar.

C. Effect of Problem Scaling

Figure 4(f) depicts the effect of migration on scaling the problem size with class B and C inputs on 16 nodes. For BT, FT and SP, we observe an increase in overhead as the task size increases from class B to class C. This behavior is expected as problem scaling results in larger data per node. However, the inverse behavior is observed for CG and LU. Though the actual downtime becomes longer as the task size increases from class B to class C, the absolute downtime (0.27-1.19 seconds) is so small that its effect is subsumed by the variance of overall job execution time (up to 11 seconds for CG and up to 8 seconds for LU), ultimately resulting in an overall decrease in overhead for increasing task sizes.

D. Effect of Task Scaling

We next examined migration under strong scaling by increasing the number of nodes. Figure 4(g) depicts the overhead for NPB codes with class C inputs on varying number of nodes (4, 8 or 9 and 16). In most cases, overheads tend to decrease as the number of nodes increases. Yet, BT(Stop&Copy), CG(Live), LU(Stop&Copy)

and SP(Live) show no obvious trends. As with problem scaling, this can be attributed to the relatively minor migration downtime, which is effectively subsumed by variances in job execution times. Hence, no measurable effect on task scaling is observed in these cases.

Figure 4(h) depicts the speedup on 4, 8 or 9 and 16 nodes normalized to the wall-clock time on 4 nodes. The figure also shows the relative speedup of live migration, of frozen migration (stop©-only) and without migration. The lightly colored portion of the bars represents the execution time of the benchmarks in the presence of one migration. The aggregate value of light and dark stacked bars presents the execution time without migration. Hence, the dark portions of the bars show the loss in speedup due to migration. The largest loss in speedup is 0.21 with FT on 16 nodes. This can be attributed to FT's relatively large migration overhead (8.5 seconds) compared to its rather short execution time (150 seconds). While the overhead increases proportionately to the memory footprint, the memory footprint is limited by system hardware (total memory size), which also limits the migration overhead. Hence, our results indicate an increasing potential for strong scaling of the benchmarks.

E. Page Access Pattern & Iterative Live Migration

Next, page access patterns of the NPB application are analyzed with respect to their effect on live migration, specifically on precopy duration and downtime. Figures 4(i) and 4(j) depict the page access pattern of CG and FT with class C inputs on 16 nodes. We sampled the number of memory pages modified per 0.1 second interval. The page size is 4KB on our experimental system. The write frequency to memory pages of CG is lower ($\sim 393/\text{interval}$) than that of FT (between $1000/\text{interval}$ and $5000/\text{interval}$). Both codes show a certain regularity in their write patterns, albeit quite an individual one each.

To compare the iterative live migration with the page access pattern, we further evaluated the aggregate amount of the transferred memory and the duration of each iteration of a full memory sweep, as depicted in Figures 4(k) and 4(l). During the first iteration, all non-empty pages are transferred (depicted as dark gray in the figures). In subsequent iterations, only those pages modified since the last iteration (light gray) are transferred. We observed that the write patterns of Figures 4(i) and 4(j) are in accord with those depicted in Figures 4(k) and 4(l). For CG, memory pages written in the last two iterations account for 1.6MB corresponding to a write frequency of 393 pages. However, the dirtied memory in the second iteration is only 1.6MB while that in the first iteration is 62.2MB overall, even though the iteration durations are the same. This happens because the first iteration occasionally coincides with the initialization phase of the application. There, higher per-page memory access frequencies (2000-8000 per 0.1 sec-

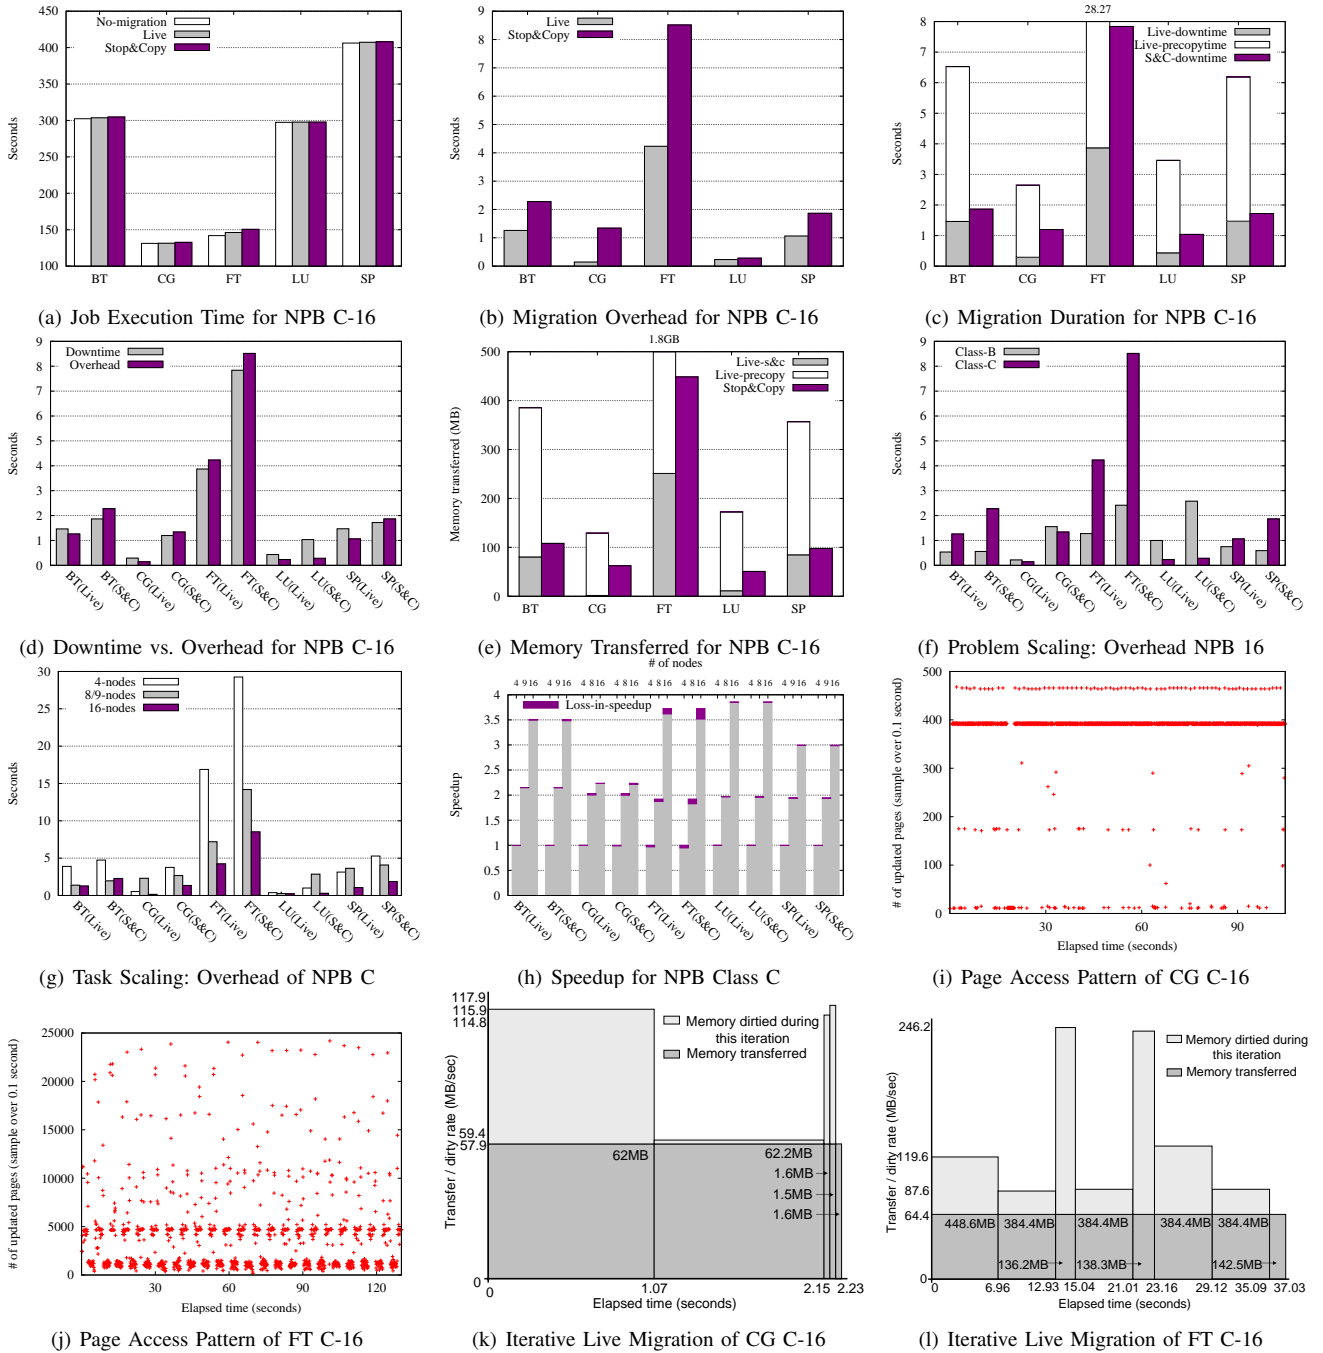


Fig. 4: Evaluation with NPB (C-16: Class C on 16 Nodes)

ond) persist combined with more frequent cold misses in cache. (This effect was experimentally confirmed but is not explicitly reflected in Figure 4(k)). For FT, 138MB and 384MB overall of dirtied memory was transferred in an alternating pattern during consecutive iterations (Figure 4(l)) corresponding to the most significant clusters at 1200 and 4600 page updates in Figure 4(j).

These profile patterns of the page modification rate could be archived and subsequently utilized as an empirical

factor for future application runs/precopy decisions to choose a more sensitive termination condition (see Section II). Besides the page access pattern, the communication pattern also affects the precopy decision. For example, if the application is both communication intensive and memory intensive, the opportunity for the precopy operation to complete before an imminent failure is low for high migration overhead and significant network contention. Three main criteria for trading off live and frozen migration and

for precopy termination conditions are:

- thresholds, *e.g.*, temperature watermarks in Section III-A, memory/difference/overhead thresholds in Section III-B1;
- available network bandwidth determined by dynamic monitoring; and
- size of the write set. (If the memory dirtying rate is faster than the available network rate, the precopy operation may be ineffective.)

Based on these conditions, a heuristics algorithm can be designed. However, the applications we evaluated are not sufficient to design such an algorithm. In future work, we plan to create and assess applications with varying communication rate and memory access pattern that are more suitable.

F. Process-Level Live Migration vs. Xen Virtualization Live Migration

We next provide the performance comparison of our approach to another solution at the OS virtualization layer in the context of proactive FT of MPI applications [33]. The common benchmarks measured with both solutions on the same hardware were NPB BT, CG, LU and SP. For these NPB codes with class C inputs on 16 nodes, the overhead of migrating an entire guest OS ranged between 4-14 seconds under Xen virtualization while the process-level solution caused only around 1 second overhead. The time taken from initiating migration to actual completion on 16 nodes ranged between 14-24 seconds for live migration as opposed to a near-constant cost of 13-14 seconds for frozen migration (stop©), both under Xen virtualization. In contrast, our process-level solution only requires 2.6-6.5 seconds for live migration and 1-1.9 seconds for frozen (stop©-only) migration. The main difference between the two solutions is that the Xen virtualization solution induced a 13 second minimum overhead to transfer the entire memory image of the inactive guest VM (with a guest memory cap of 1GB) while the process-level solution constrained migration to only the memory of the respective process. Hence, our solution (with 1-6.5 seconds of prior warning to successfully trigger live process migration) significantly outperforms the Xen virtualization solution (with 13-24 seconds of prior warning). One could argue that only a subset of the OS image needs to be migrated, yet the strength of virtualization lies in its transparency, yet it comes at the cost of indiscriminate transfer of the entire virtual memory range.

G. Proactive FT Complements Reactive FT

We claim that proactive FT complements its reactive sister. This was already empirically shown in Figure 4(h) where we noted that scalability of live migration depends on the amount of transferred memory. Once local memory is exhausted by an application, the overhead of a single migration will remain constant irrespective of the number

of nodes. Of course, the rate of failures in larger systems is bound to increase, but proactive FT supports larger systems while reactive schemes result in increased I/O bandwidth requirements, which can become a bottleneck. This supports our argument that proactive schemes are important since they can complement reactive schemes in lowering checkpoint frequency requirements of the latter.

An analytical argument for the complementary nature is given next. The objective here is to assess the ability of proactive FT to successfully counter imminent faults, which subsequently allows reactive schemes to engage in less frequent checkpointing. Let the time interval between checkpoints be T_c , the time to save checkpoint information be T_s , and the mean time between failures (MTBF) be T_f . Then, the optimal checkpoint rate is $T_c = \sqrt{2 \times T_s \times T_f}$ [53]. We also observed that the mean checkpoint time (T_s) for BT, CG, FT, LU and SP with class C inputs on 4, 8 or 9 and 16 nodes is 23 seconds on the same experimental cluster [51]. With a MTBF of 1.25 hours [35], the optimal checkpoint rate T_c is $T_c = \sqrt{2 \times 23 \times (1.25 \times 60 \times 60)} = 455$ seconds. Let us further assume that 70% of failures can be predicted [40] and can thus be avoided by our proactive migration. (Sahoo *et al.* actually reported that up to 70% of failures can be predicted without prior warning; with a prior warning window, the number of proactively tolerated failures could even exceed 70%.) Our solution can then prolong reactive checkpoint intervals to $T_c = \sqrt{2 \times 23 \times (1.25 / (1 - 0.7) \times 60 \times 60)} = 831$ seconds. The challenge with proactive FT then becomes (1) to provide a sufficient number of spare nodes to avoid initial failures by live or frozen migration and (2) to repair faulty nodes in the background such that jobs running over the course of days can reclaim failed nodes for future proactive FT.

VI. Related Work

Process-level migration has been studied extensively in the past [32], [37], [46], [26], [4], [11], [15]. In the context of HPC, many MPI implementations have been retrofitted with or design for FT, ranging from automatic methods (checkpoint-based or log-based) [44], [41], [5] to non-automated approaches [3], [17].

System checkpoint-based methods commonly employ a combination of OS support to checkpoint a process image (*e.g.*, via the BLCR Linux module [15]) combined with a coordinated checkpoint negotiation using collective communication among MPI tasks. User-level checkpointing, in contrast, relies on runtime library support and may require source preprocessing of an application and typically inflicts lower overhead, yet fails to capture critical system resources, such as file descriptors [28], [36]. Log-based methods exploit messages logging and optionally their temporal ordering, where the latter is required for

asynchronous non-coordinated checkpointing. MPICH-V [5] implements three such protocols. It uses Condor’s user-level checkpoint library [29]. Non-automatic approaches generally involve explicit invocation of checkpoint routines. In contrast to these methods, ghost process [49] provides a global management service with process virtualization mechanisms supporting process migration within clusters within the Karrighed single-image system.

Recent studies focus on proactive FT. The feasibility of proactive FT has been demonstrated at the job scheduling level [34] and in Adaptive MPI [8], [7], [9] using a combination of (a) object virtualization techniques to migrate tasks and (b) causal message logging [16] within the MPI runtime system of Charm++ applications. In contrast to Charm++, our solution is coarser grained as FT is provided at the process level, thereby encapsulating most of the process context, including open file descriptors, which are beyond the MPI runtime layer.

Two facets on proactive FT are intensively studied. First, there are a number of research efforts on failure prediction [40], [20], [21]. These papers report high failure prediction accuracy with a prior warning window, which is the premise for our process migration mechanism proposed in this paper. Second, various migration mechanisms are proposed as discussed previously. Furthermore, MPI-Mitten [12], an MPI library between the MPI layer and the application layer, provides proactive fault tolerance to MPI applications. It uses HPCM [13] as a middleware to support user-level heterogeneous process migration. These two facets are integrated in approaches that combine prediction and migration in proactive FT systems and evaluate different FT policies. In [50], the authors provide a generic framework based on a modular architecture allowing the implementation of new proactive fault tolerance policies/mechanisms. An agent oriented framework [23] was developed for grid computing environments with separate agents to monitor individual classes or subclasses of faults and proactively act to avoid or tolerate a fault. Sun *et al.* provide fault-aware systems, such as FARS [27] and FENCE [45], to increase the accuracy of fault prediction and improve system resilience to failures with different fault management mechanisms including process migration. They also model the migration cost and introduce a dynamic scheduling mechanism accordingly [14]. In their paper, Tikotekar *et al.* also present a simulation framework that evaluates different FT mechanisms and policies, including a combination of reactive FT and proactive FT to decrease the number of checkpoints [48], which obtained the best results among all the real and simulated FT mechanisms and policies. These prior works with their fault models, FT mechanisms for fault occurrences and their evaluation simulations, confirm that the process migration is a suitable approach for proactive FT with

lower cost than OS virtualization, which reinforces the significance of our solution.

Our work enhances LAM/MPI and BLCR [43], [15], [41], which previously was restricted to reactive FT, to a proactive live migration scheme. LAM/MPI+BLCR originally required a complete system restart for roll-back to the last checkpoint upon failure, but a number of approaches have been designed to allow (a) selected checkpoint images to be restarted on new nodes [6], (b) node and head-node failure [47], and (c) a job-pause mechanism that supports migration without restart [51]. Our solution orchestrates the BMC/IPMI health monitoring, the new fundamental utilities provided at BLCR and the extended communication mechanism at LAM/MPI through the decentralized scheduler. The framework is simple and applicable to arbitrary MPI implementations in HPC environment. Furthermore, our approach provides a *live migration mechanism*, which supports continued execution of MPI applications *during* much of the migration time. This solution parallels live migration at the OS virtualization layer [10], which has been studied in the context of proactive FT of MPI applications [33], an approach that supports integrated health-based monitoring and proactive live migration over Xen guests. We contribute *process-level* live migration and demonstrate its superior efficiency to of OS-level virtualization. In HPC, process-level solutions are more widely accepted than OS virtualization, not the least because of potential performance penalties of network virtualization or additional driver development for virtualization-bypass technologies [31], [30].

VII. Conclusion

We have contributed a novel *process-level live migration* mechanism with a concrete implementation as an enhancement to the Linux BLCR module and an integration within LAM/MPI. By monitoring the health of each node, a process can be migrated away from nodes subject to imminent failure. We show live migrations to be beneficial to frozen migration due to a lower overall overhead on application wall-clock execution time. This lower overhead is attributed to the asynchronous transfer of a large portion of a migrated process image while the application stays alive by proceeding with its execution. Process-level migration is also shown to be significantly less expensive than migrating entire OS images under Xen virtualization. The resulting proactive approach complements reactive checkpoint/restart schemes by avoiding roll-backs if failures are predicted only seconds before nodes cease to operate. Thus, the optimal number of checkpoints of applications can be nearly cut in half when 70% of failures are addressed proactively.

References

- [1] “Advanced configuration & power interface,” <http://www.acpi.info>.

- [2] "Readable dirty-bits for IA64 linux," <https://www.gelato.unsw.edu.au/archives/gelato-technical/2005-November/001080.html>.
- [3] R. T. Aulwes, D. J. Daniel, N. N. Desai, R. L. Graham, L. D. Risinger, M. A. Taylor, T. Woodall, and M. Sukalski, "Architecture of LA-MPI, a network-fault-tolerant MPI," in *IPDPS*, 2004.
- [4] A. Barak and R. Wheeler, "MOSIX: An integrated multiprocessor UNIX," in *Proceedings of the Winter 1989 USENIX Conference*. Berkeley, CA, USA: USENIX, 1989, pp. 101–112.
- [5] G. Bosilca, A. Boutellier, and F. Cappello, "MPICH-V: Toward a scalable fault tolerant MPI for volatile nodes," in *Supercomputing*, Nov. 2002.
- [6] J. Cao, Y. Li, and M. Guo, "Process migration for MPI applications based on coordinated checkpoint," in *ICPADS*, 2005, pp. 306–312.
- [7] S. Chakravorty, C. Mendes, and L. Kale, "Proactive fault tolerance in MPI applications via task migration," in *HiPC*, 2006.
- [8] S. Chakravorty, C. Mendes, and L. Kale, "Proactive fault tolerance in large systems," in *HPCRI: 1st Workshop on High Performance Computing Reliability Issues*, in *Proceedings of HPCA-11*, 2005.
- [9] S. Chakravorty, C. Mendes, and L. Kale, "A fault tolerance protocol with fast fault recovery," in *IPDPS*, 2007.
- [10] C. Clark, K. Fraser, S. Hand, J. Hansem, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *NSDI*, May 2005.
- [11] F. Douglass and J. K. Ousterhout, "Transparent process migration: Design alternatives and the sprite implementation." *Softw., Pract. Exper.*, vol. 21, no. 8, pp. 757–785, 1991.
- [12] C. Du and X.-H. Sun, "MPI-Mitten: Enabling migration technology in MPI," in *IEEE CCGrid*, 2006.
- [13] C. Du, X.-H. Sun, and K. Chanchio, "HPCM: A pre-compiler aided middleware for the mobility of legacy code," in *IEEE Cluster*, 2003.
- [14] C. Du, X.-H. Sun, and M. Wu, "Dynamic scheduling with process migration," in *IEEE CCGrid*, May 2007.
- [15] J. Duell, "The design and implementation of berkeley lab's linux checkpoint/restart," Lawrence Berkeley National Laboratory, TR, 2000.
- [16] E. N. Elnozahy and W. Zwaenepoel, "Manetho: Transparent roll back-recovery with low overhead, limited rollback, and fast output commit," *IEEE Trans. Comput.*, vol. 41, no. 5, pp. 526–531, 1992.
- [17] G. E. Fagg and J. J. Dongarra, "FT-MPI: Fault Tolerant MPI, supporting dynamic applications in a dynamic world," in *Euro PVM/MPI User's Group Meeting*, vol. 1908, 2000, pp. 346–353.
- [18] S. Ghemawat, H. Gobioff, and S.-T. Leung, "The Google file system," in *SOSP '03*, 2003, pp. 29–43.
- [19] R. Gioiosa, J. C. Sancho, S. Jiang, and F. Petrini, "Transparent, incremental checkpointing at kernel level: a foundation for fault tolerance for parallel computers," in *Supercomputing*, 2005.
- [20] X. Gu, S. Papadimitriou, P. S. Yu, and S.-P. Chang, "Toward predictive failure management for distributed stream processing systems," in *IEEE ICDCS*, 2008.
- [21] P. Gujrati, Y. Li, Z. Lan, R. Thakur, and J. White, "A meta-learning failure predictor for BlueGene/L systems," in *ICPP*, Sep. 2007.
- [22] C.-H. Hsu and W.-C. Feng, "A power-aware run-time system for high-performance computing," in *Supercomputing*, 2005.
- [23] M. T. Huda, H. W. Schmidt, and I. D. Peake, "An agent oriented proactive fault-tolerant framework for grid computing," in *International Conference on e-Science and Grid Computing*, 2005.
- [24] J. Hursey, J. M. Squyres, and A. Lumsdaine, "A checkpoint and restart service specification for Open MPI," Indiana University, Computer Science Department, Technical Report, 2006.
- [25] J. Hursey, J. M. Squyres, T. I. Mattox, and A. Lumsdaine, "The design and implementation of checkpoint/restart process fault tolerance for Open MPI," in *DPDNS*, Mar. 2007.
- [26] E. Jul, H. M. Levy, N. C. Hutchinson, and A. P. Black, "Fine-grained mobility in the emerald system." *ACM Trans. Comput. Syst.*, vol. 6, no. 1, pp. 109–133, 1988.
- [27] Y. Li, P. Gujrati, Z. Lan, and X.-H. Sun, "Fault-driven re-scheduling for improving system-level fault resilience," in *ICPP*, 2007.
- [28] M. Litzkow, "Remote unix - turning idle workstations into cycle servers," in *Usenix Summer Conference*, 1987, pp. 381–384.
- [29] M. Litzkow, T. Tannenbaum, J. Basney, and M. Livny, "Checkpoint and migration of UNIX processes in the Condor distributed processing system," University of Wisconsin - Madison Computer Sciences Department, Tech. Rep. UW-CS-TR-1346, April 1997.
- [30] J. Liu, W. Huang, B. Abali, and D. Panda, "High performance vmm-bypass I/O in virtual machines," in *USENIX Conference*, Jun. 2006.
- [31] A. Menon, A. Cox, and W. Zwaenepoel, "Optimizing network virtualization in Xen," in *USENIX Conference*, Jun. 2006.
- [32] D. S. Milojevic, F. Douglass, Y. Paindaveine, R. Wheeler, and S. Zhou, "Process migration," *ACM Computing Surveys (CSUR)*, vol. 32, no. 3, pp. 241–299, 2000.
- [33] A. B. Nagarajan and F. Mueller, "Proactive fault tolerance for HPC with Xen virtualization," in *ICS*, Jun. 2007.
- [34] A. Oliner, R. Sahoo, J. Moreira, M. Gupta, and A. Sivasubramaniam, "Fault-aware job scheduling for BlueGene/L systems," in *IPDPS*, 2004.
- [35] I. Philp, "Software failures and the road to a petaflop machine," in *HPCRI: 1st Workshop on High Performance Computing Reliability Issues*, in *Proceedings of HPCA-11*. IEEE Computer Society, 2005.
- [36] J. S. Plank, M. Beck, G. Kingsley, and K. Li, "Libckpt: Transparent checkpointing under Unix," in *Usenix Winter Technical Conference*, January 1995, pp. 213–223.
- [37] M. L. Powell and B. P. Miller, "Process migration in DEMOS/MP," in *Symposium on Operating Systems Principles*, Oct. 1983.
- [38] H. Quinn and P. Graham, "Terrestrial-based radiation upsets: A cautionary tale," in *FCCM '05*, Apr. 18–20, 2005.
- [39] S. Rani, C. Leangsuksun, A. Tikotekar, V. Rampure, and S. Scott, "Toward efficient failure detection and recovery in HPC," in *High Availability and Performance Computing Workshop*, 2006.
- [40] R. Sahoo, A. Oliner, I. Rish, M. Gupta, J. Moreira, S. Ma, R. Vilalta, and A. Sivasubramaniam, "Critical event prediction for proactive management in large-scale computer clusters," in *KDD '03*, 2003.
- [41] S. Sankaran, J. M. Squyres, B. Barrett, A. Lumsdaine, J. Duell, P. Hargrove, and E. Roman, "The LAM/MPI checkpoint/restart framework: System-initiated checkpointing," in *LACSI*, Oct. 2003.
- [42] H. Song, C. Leangsuksun, and R. Nassar, "Availability modeling and analysis on high performance cluster computing systems," in *ARES*, 2006, pp. 305–313.
- [43] J. M. Squyres and A. Lumsdaine, "A Component Architecture for LAM/MPI," in *European PVM/MPI Users' Group Meeting*, ser. Lecture Notes in Computer Science, no. 2840. Venice, Italy: Springer-Verlag, September / October 2003, pp. 379–387.
- [44] G. Stellner, "CoCheck: checkpointing and process migration for MPI," in *Proceedings of IPPS '96*, 1996.
- [45] X.-H. Sun, Z. Lan, Y. Li, H. Jin, and Z. Zheng, "Towards a fault-aware computing environment," in *HAPCV*, Mar. 2008.
- [46] M. Theimer, K. A. Lantz, and D. R. Cheriton, "Preemptable remote execution facilities for the V-System," in *SOSP*, 1985, pp. 2–12.
- [47] A. Tikotekar, C. Leangsuksun, and S. L. Scott, "On the survivability of standard MPI applications," in *LCI International Conference on Linux Clusters: The HPC Revolution*, May 2006.
- [48] A. Tikotekar, G. Vallée, T. Naughton, S. L. Scott, and C. Leangsuksun, "Evaluation of fault-tolerant policies using simulation," in *IEEE Cluster*, Sep. 17–20, 2007.
- [49] G. Vallee, R. Lottiaux, D. Margery, C. Morin, and J.-Y. Berthou, "Ghost process: a sound basis to implement process duplication, migration and checkpoint/restart in linux clusters," *ISPDC*, 2005.
- [50] G. Vallée, K. Charoenpornwattana, C. Engelmann, A. Tikotekar, C. B. Leangsuksun, T. Naughton, and S. L. Scott, "A framework for proactive fault tolerance," in *ARES*, 2007, pp. 659–664.
- [51] C. Wang, F. Mueller, C. Engelmann, and S. Scott, "A job pause service under LAM/MPI+BLCR for transparent fault tolerance," in *IPDPS*, Apr. 2007.
- [52] F. Wong, R. Martin, R. Arpaci-Dusseau, and D. Culler, "Architectural requirements and scalability of the NAS parallel benchmarks," in *Supercomputing*, 1999.
- [53] J. W. Young, "A first order approximation to the optimum checkpoint interval," *Commun. ACM*, vol. 17, no. 9, pp. 530–531, 1974.