



2023 OFA Virtual Workshop


# **/CHECK: LEVERAGING RDMA AND MALLEABILITY FOR APPLICATION-LEVEL CHECKPOINTING IN HPC SYSTEMS**

Jophin John

Chair of Computer Architecture and Parallel Systems (CAPS)  
Technical University of Munich, Germany



# OVERVIEW



# A CASE FOR DYNAMISM IN CHECKPOINTING

- Plenty of techniques for fault tolerance
- Focus on Application-level checkpoint restart (widely used in Simulations)
- Systems and Applications are becoming malleable in HPC
- *Dynamism* in checkpointing can improve
  - Application performance
  - System utilisation
- RDMA can be used for efficient checkpoint management
- *iCheck* – a fully adaptive [Invasive Checkpoint Management System](#)



# INVASIVE COMPUTING



4



© OpenFabrics Alliance

# INVASIVE COMPUTING

- DFG Transregio Research Centre TRR89 "Invasive Computing" – InvasIC



- **Focus**

- Dynamic resource management on massively parallel chip multiprocessors
- Resource-aware applications: invading, infecting, retreating from resources
- Integration: Hardware - OS – Language & Compiler - Tools – Applications



- **Investigation in the context of HPC – Technical University of Munich**

- Chair for Scientific Computing: Hans Bungartz, Michael Bader
- Chair for Computer Architecture and Parallel Systems: Michael Gerndt



# **/CHECK & MALLEABILITY**

- **iCheck supports invasive/malleable applications developed using Invasive (Malleable) HPC Infrastructure**
- **Invasive HPC Infrastructure**
  - System level
    - Malleable Resource and Job Management System – ***iRM***
    - Resource aware MPI – ***iMPI***
  - Application level
    - Programming models – Elastic Phase Oriented Programming model (EPOP)
    - Applications – Tsunami Simulation, ***iMD***, ***iHeat***, ***iSWE***, ***iSum***
  - Services
    - Data analytics application support (Using Apache Spark)
    - Power budget enforcement
    - Fault tolerance - ***iCheck***



# INVASIVE INFRASTRUCTURE + ICHECK

## Invasive Resource Manager (*iRM*)

- Extension of SLURM with dynamic resource management

## Invasive Message Passing Interface (*iMPI*)


- Extension of MPICH
- Four new operations for dynamic processes
  - *MPI\_Init\_adapt*
  - *MPI\_Probe\_adapt*
  - *MPI\_Comm\_adapt\_begin*
  - *MPI\_Comm\_adapt\_commit*

I. A. Comprés Ureña and Michael Gerndt. Towards Elastic Resource Management. In *Tools for High Performance Computing 2017*, pages 105–127. Springer International Publishing, 2019

Ao Mo-Hellenbrand. Resource-Aware and Elastic Parallel Software Development for Distributed-Memory HPC Systems. Dissertation, Technische Universität München, 2019


# INVASIVE RESOURCE MANAGER - iRM

- iRM decides about resource reconfiguration
- Application expansion or reduction is possible
- Application reacts to the resource changes triggered by iRM
- Resource change is a six-step process




# INVASIVE RESOURCE MANAGER - iRM

- iRM decides about resource reconfiguration
- Application expansion or reduction is possible
- Application reacts to the resource changes triggered by iRM
- Resource change is a six-step process
- Scheduler Plugin is created for iCheck
- iCheck interacts with the scheduler



# INVASIVE MPI - /MPI



## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)
        // Redistribute data
        MPI_Comm_adapt_commit ( );
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt (...)
            if resource_change {
                MPI_Comm_adapt_begin (...)
                // Redistribute data
                MPI_Comm_adapt_commit ( );
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
    ...
}
```

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)
        // Redistribute data
        MPI_Comm_adapt_commit ( );
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt (...)
            if resource_change {
                MPI_Comm_adapt_begin (...)
                // Redistribute data
                MPI_Comm_adapt_commit ( );
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ...
    }
    // End elastic block n
    phase index++;
    ...
}
```

Flow of initial set of processes

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)
        // Redistribute data
        MPI_Comm_adapt_commit ( );
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt (...)
            if resource_change {
                MPI_Comm_adapt_begin (...)
                // Redistribute data
                MPI_Comm_adapt_commit ( );
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
    ...
}
```

Flow of newly added set of processes

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)

        // Redistribute data
        MPI_Comm_adapt_commit ( );
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt (...)

            if resource_change {
                MPI_Comm_adapt_begin (...)

                // Redistribute data
                MPI_Comm_adapt_commit ( );
            }

            iteration number++;

            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
    ...
}
```

Flow of initial set of processes

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)

        // Redistribute data
        MPI_Comm_adapt_commit ( );
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt (...)

            if resource_change {
                MPI_Comm_adapt_begin (...)

                // Redistribute data
                MPI_Comm_adapt_commit ( );
            }

            iteration number++;

            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
}
```

Flow of initial set of processes

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin ...
        // Redistribute data
        MPI_Comm_adapt_commit ();
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt ...
            if resource_change {
                MPI_Comm_adapt_begin ...
                // Redistribute data
                MPI_Comm_adapt_commit ();
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
    ...
}
```

Joining processes

Existing processes

# Pseudocode Invasive MPI+/CHECK

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin (...)  

icheck_redistribute();
        MPI_Comm_adapt_commit () ;
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while ( block_condition ) {
            MPI_Probe_adapt ...
            if resource_change {
                MPI_Comm_adapt_begin (...)  

icheck_redistribute();
                MPI_Comm_adapt_commit () ;
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ...
    }
    // End elastic block n
    phase index++;
    ...
}
```

Joining processes

Existing processes

## Pseudocode - a simple Invasive MPI Application

```
int main() {
    ...
    MPI_Init_adapt(...,mytype)
    // Initialization block
    if mytype == initial_process {
        set phase index = 0
    }
    else {
        // Newly joining processes
        MPI_Comm_adapt_begin ...
        // Redistribute data
        MPI_Comm_adapt_commit();
    }
    // Begin elastic block 1
    if (phase index == 0) {
        while (block_condition) {
            MPI_Probe_adapt ...
            if resource_change {
                MPI_Comm_adapt_begin ...
                // Redistribute data
                MPI_Comm_adapt_commit();
            }
            iteration number++;
            // Compute Intensive part
        }
    }
    // End elastic block 1
    ...
    // Begin elastic block n
    if (phase index == n){ ... }
    // End elastic block n
    phase index++;
    ...
}
```


All processes starts working together



# /CHECK



# iCHECK ARCHITECTURE

- iCheck – An invasive RDMA based multilevel application level checkpointing system
- Deployed in dedicated nodes
- iCheck Core
  - Controller – *single instance component*
  - Manager – *one instance per iCheck node*
  - Agents – *multiple instances per iCheck node*
- iCheck Library
  - Application interacts with iCheck Core
- Simultaneous checkpoint management of multiple applications
  - Potential offered by dynamism is enormous




# iCHECK WORKFLOW

- Application registers with the controller
- Controller assigns agents and iCheck nodes
  - Agent placement algorithm
  - Node selection algorithm
- Manager launches agents
- Agents connect with the applications
- RDMA configuration performed
- Application calls commit & continues the execution
  - Agent retrieves the checkpoint using RMA
- Application can probe for agent change



# RDMA IN iCHECK

- Application registers memory regions
- Uses Libfabric library
- Multiple techniques in *iCheck*
  - RDMA only using memory regions in libfabric
  - RDMA + Shared memory
- Push and Pull Techniques
  - Push: Application pushes checkpoint data to agents
  - Pull: Agents read checkpoint data from applications
- Agents write data to PFS






# DYNAMISM IN ICHECK

# DYNAMISM IN iCHECK


- System-level dynamism in *iCheck*
  - Scaling of *iCheck nodes* (Manager) using *iRM*
  - *Agent* behaviour
- Application-level dynamism
  - Scaling of *Agents*
  - Support for *Malleability*



Hierarchical view of *iCheck* system


# SYSTEM-LEVEL DYNAMISM – *iCHECK NODES*

- Created a new Slurm plugin to support *iCheck*
- iRM can reconfigure *iCheck* nodes (Naive approach)
  - *iCheck* can also request for new nodes




# SYSTEM-LEVEL DYNAMISM – *ICHECK NODES*

- Created a new Slurm plugin to support *iCheck*
- iRM can reconfigure *iCheck* nodes (Naive approach)
  - *iCheck* can also request for new nodes
- Complex strategies are necessary




# SYSTEM-LEVEL DYNAMISM - AGENTS

- Agents are lowest level component in *iCheck* system
- Agents write/read checkpoint from applications
- Agents can do more than checkpoint retrieval




# SYSTEM-LEVEL DYNAMISM - AGENTS

- Agents are lowest level component in *iCheck* system
- Agents write/read checkpoint from applications
- Agents can do more than checkpoint retrieval
- *iCheck* offers different plugins
  - New plugins can be added as requirement and technology changes
  - Does not need to worry about optimisation




# SYSTEM-LEVEL DYNAMISM - AGENTS

- Agents are lowest level component in *iCheck* system
- Agents write/read checkpoint from applications
- Agents can do more than checkpoint retrieval
- *iCheck* offers different plugins
  - New plugins can be added as requirement and technology changes
  - Does not need to worry about optimisation
- Checkpointing systems can do a lot more on-the-fly




# APPLICATION-LEVEL DYNAMISM - AGENTS

- Basic API
  - `icheck_init("appname", ..., status);`
  - `icheck_add("var_name", &var, SIZE);`
  - `icheck_commit();`
  - `icheck_restart();`
  - `icheck_restore(var_name,&var)`
  - `icheck_finalize ( icheck state );`
  - `icheck_enable_async()` – enable asynchronous checkpointing
- API call `icheck_probe_agents()` will contact controller
- RDMA reconfiguration
- Controller can decide on agent change



# APPLICATION-LEVEL DYNAMISM - AGENTS

- Basic API
  - `icheck_init("appname", ..., status);`
  - `icheck_add("var_name", &var, SIZE);`
  - `icheck_commit();`
  - `icheck_restart();`
  - `icheck_restore(var_name,&var)`
  - `icheck_finalize ( icheck state );`
  - `icheck_enable_async()` – enable asynchronous checkpointing
- API call `icheck_probe_agents()` will contact controller
- Controller can decide on agent change
- RDMA reconfiguration
- Performs agent change on-the-fly based on the application metrics



# APPLICATION-LEVEL DYNAMISM - *MALLEABILITY*

- Support checkpoint restart in malleable applications
- Perform data redistribution for malleable applications
- Malleable API
  - *icheck\_add\_adapt*("var\_name", &var, sizeof(int), ..., **TYPE**);
  - *icheck\_redistribute*(var\_name,&var,..., **TYPE**)

Pseudocode using iCheck API

```
1 #include<icheck.h>
2 int main() {
3     MPI_Init_adapt(..., type)
4     float data[SIZE];
5     icheck_init(..., type);
6     icheck_add_adapt("data", data, ..., BLOCK);
7     if(checkpoint_available && no_adapt){
8         icheck_restart();
9     }
10    if (type == joining) {
11        MPI_Comm_adapt_begin(...);
12        icheck_redistribute("data", data,
13                             new_size, BLOCK)
14        MPI_Comm_adapt_commit();
15    }
16    while (true){
17        MPI_Probe_adapt(resource_change ,...);
18        if (resource_change) {
19            MPI_Comm_adapt_begin(...);
20            icheck_redistribute("data", data,
21                                new_size, BLOCK)
22            MPI_Comm_adapt_commit();
23        }
24        /*Read/Write data[] */
25        if(iteration %100)
26            icheck_commit();
27        /*Check for agent change*/
28        if(iteration %1000)
29            icheck_probe_agents();
30    }
31    icheck_finalize(IC_PERSIST);
32    MPI_Finalize();
33 }
```



# EVALUATION

# EVALUATION SETUP



- Tests were performed on SuperMUC-NG
- Evaluation on 16 nodes (768 cores)
  - 12 nodes (576 cores) for applications
  - Four nodes for iCheck
- Four applications
  - ls1 mardyn
  - LULESH
  - 2D Jacobi heat simulation
  - Synthetic application – checkpoints 2.3GB per second - 1million floats
- Scenarios:
  - Compare iCheck vs MPI-IO in ls1 mardyn
  - Compare iCheck vs MPI-IO vs SCR in a synthetic application
  - Effect of dynamic agents on a 2D heat simulation application
  - Effect of different agent placement strategies using multiple synthetic applications



SuperMUC-NG at LRZ, Germany  
<https://doku.lrz.de/display/PUBLIC/SuperMUC-NG>

# RESULTS

## iCheck vs MPI-IO in ls1 mardyn




*Up to 5000 times (best case with 12 agents) faster checkpointing with iCheck.*

*In the worst case, iCheck is 400 times faster (single agent).*

# RESULTS

## iCheck vs MPI-IO vs SCR in Synthetic application




*Up to 100x and 57x faster checkpointing (using 12 agents) than MPI-IO and SCR*

*In the worst case (using single agent), iCheck is 8x faster than MPI-IO and 4x faster than SCR*

# RESULTS

## Effect of dynamic agents on a 2D heat simulation application




- Dynamism might not be ideal in all cases
- Heuristics are needed to model the performance characteristics

Up to 50 times faster performance possible on the fly

# RESULTS


## Effect of different agent placement strategies in synthetic applications



- Stacking: all agents in same iCheck node
- Bandwidth: placed agents based on available bandwidth
- CP Interval: placed agents on checkpoint interval

Significant improvement in bandwidth utilisation with ideal agent placement

# CHALLENGES



## Data redistribution

Redistribution not trivial  
Complex datatypes



## Checkpoint management

Agent placement  
Efficient resource utilisation



## Privacy and Data

Multiple applications  
IO operations

# INTO THE FUTURE

	Optimisation potential	Pluggable services Hardware
	Data & Resource management	Malleable Non-Malleable
	Deployment in Cloud	HPC on Cloud iCheck can be deployed easily
	Policies for checkpoint management (not limited to performance improvement)	Power aware Carbon aware
	Support for Accelerators	GPUS FPGA



2023 OFA Virtual Workshop

# THANK YOU

Jophin John

Chair of Computer Architecture and Parallel Systems  
Technical University of Munich, Germany



CAPS