ECOFIT: A Framework to Estimate Energy Consumption of Fault Tolerance protocols during HPC executions

M.E.M. Diouri¹, O. Glück¹, L. Lefèvre¹, and F. Cappello²

1- INRIA Avalon Team, LIP Laboratory (CNRS, ENS, INRIA, University of Lyon) 2- INRIA - University of Illinois - Argonne National Laboratory

Lyon (France), June 13th, 2013

9th JLPC Workshop

1/24

Design of the ECOFIT framework Validation of our Energy Estimating Approach Energy-Aware Protocol Selection Conclusion

Context

A wide range of scientific applications :

A new performance objective for the end of the decade: Exascale.

Several millions of CPU cores running up to a billion of threads

Will experience various kind of faults many times per day => Fault Tolerance (FT) is mandatory !

Power consumption: a limiting factor to the future growth in HPC

2/24

(ロ) (型) (ヨ) (ヨ) (ヨ) (コ)

Design of the ECOFIT framework Validation of our Energy Estimating Approach Energy-Aware Protocol Selection Conclusion

Context

However, fault tolerance and energy consumption are interrelated:

Must address both power/energy consumption and fault tolerance

3/24

イロト (得) (ヨ) (ヨ) (ヨ) ()

Power consumption of FT depends on many parameters (FTXS'2012)

The best protocol depends on the execution configuration

Design of the ECOFIT framework Validation of our Energy Estimating Approach Energy-Aware Protocol Selection Conclusion

Context

Currently, evaluate the energy of FT = measure it with wattmeters

- Measuring accurately the energy consumption is hard ! (EE-LSDS'2013)
- Not practical for protocol selection: not before the execution

An accurate estimator of the power consumption for FT protocols

- For any execution configurations.
- Compare FT protocols in a given execution configuration.

4/24

Design of the ECOFIT framework Validation of our Energy Estimating Approach Energy-Aware Protocol Selection Conclusion

Current Fault tolerance protocols

3 categories of protocols:

uncoordinated, coordinated, hierarchical protocols.

Rely on checkpointing/restart:

- with message logging in uncoordinated protocols
- with process synchronization in coordinated protocols.

5/24

(ロ) (型) (ヨ) (ヨ) (ヨ) (コ)

In hierarchical protocols: processes organized in clusters.

- process synchronization inside a same cluster.
- message logging between clusters.

Design of the ECOFIT framework Validation of our Energy Estimating Approach Energy-Aware Protocol Selection Conclusion





- 2 Design of the ECOFIT framework
- 3 Validation of our Energy Estimating Approach

6/24

◆□ ▶ ◆□ ▶ ◆ □ ▶ ◆ □ ▶ ◆ □ ● ● ● ●

4 Energy-Aware Protocol Selection

5 Conclusion

ECOFIT: High-level operations

- Checkpointing
 - storing a snapshot image of the current application state
- Message logging
 - saving on each sender process the messages sent
- Coordination
 - synchronizing the processes before taking the checkpoints
- Recovery
 - in case of failure: restarting the execution of the application from the last checkpoint.

7/24

イロト イポト イヨト イヨト

ECOFIT: Basic operations

- Checkpointing: writing a data on a media storage.
- Message logging: writing a data on a media storage.
- Coordination:
 - active polling during the transmission of inflight messages

・ロト ・得ト ・ヨト ・ヨト … ヨ

- synchronization when no more inflight message.
- Recovery:
 - restarting: reading a data from a media storage

8/24

• application re-execution (next future :-))

ECOFIT: Associated parameters

Estimating the energy consumption of a high-level operation *op* really complex: a large set of parameters

- protocols:
 - checkpointing interval, checkpointing storage destination, etc.
- application:
 - number of processes, number and size of messages exchanged, volume of data written/read by each process, etc.
- hardware:
 - number of cores per node, memory architecture, type of hard disk drives, etc.

ECOFIT integrates an automated calibration component

• in order to take into consideration all the parameters

Calibration

- Calibration: gather energy knowledge of all the basic operations according to the hardware
- A set of simple benchmarks: extract the energy consumption ξ_{op} of the basic operations in FT protocols.

10/24

・ロト ・雪 ・ ・ ヨ ・ ・ ヨ ・ うらる

- Goal = make our energy estimations accurate on any supercomputer
- It needs to be done only occasionally

Energy consumption of a node *i* during *op*: $\xi_{op}^{i} = \rho_{op}^{i} \cdot t_{op}^{j}$.

Power calibration

Power consumption of an operation op: $\rho_{op}^{i} = \rho_{idle}^{i} + \Delta \rho_{op}^{i}$

 ρ_{idle}^{i} different for identical nodes (IGCC'2013) => Measure ρ_{idle}^{i} for each node.

For a given *op*, $\Delta \rho_{op}^{i}$ is the same on identical nodes => Measure $\Delta \rho_{op}^{i}$ during each *op* for each type of nodes.

ECOFIT calibrates $\Delta \rho_{op}^{i}$ by varying the number of cores that perform the same *op*.

• In order to take into account the impact of parallelism

11/24

・ロト ・日 ・ ・ ヨ ・ ・ ヨ ・ うらう

Execution time calibration

Checkpointing, message logging or restarting

$$t_{op}^{i} = t_{access}^{i} + t_{transfer}^{i} = t_{access}^{i} + \frac{V_{data}}{r_{transfer}^{i}}$$

A simple benchmark that measures t_{op}^{i} for different values of V_{data}

Parallelism: Calibration for different numbers of processes per node

For all the different storage medium (RAM, HDD, SSD, NFS, ...)

12/24

・ロト ・日 ・ ・ ヨ ・ ・ ヨ ・ うらう

Estimation Framework



Calibration performed when a change occurs in the hardware used

13/24

Estimation Approach: Message Logging

From the user: total number and size of sent messages, number of nodes and number of processes per node.

• computes V_{data}^{mean} sent (so logged) by each node.

From the calibrator: $t_{logging}^{i}$ corresponding to V_{data}^{mean} for each node and according to the number of processes per node.

- If V_{data}^{mean} is not a size recorded by the calibrator
 - computes equation that gives $t_{logging}^{i}$ according to V_{data}
 - adjusts the equation using the method of least squares.

14/24

Experimental infrastructure

Calibrate and run HPC applications on a cluster of Grid5000

- 16 identical nodes Dell R720
- 2 Intel Xeon CPU 2.3 GHz, with 6 cores each
- 32 GB of memory
- a 10 Gigabit Ethernet network
- a HDD with a storage capacity of 598 GB

external power meters from the SME Omegawatt.

• mean power consumption, at each second for each node

15/24

◆□ ▶ ◆□ ▶ ◆ □ ▶ ◆ □ ▶ ◆ □ ◆ ○ ●

Compute the mean value over 30 experiments.

Power calibration

Measure ρ^i_{idle} of each node iCalibrate $\Delta \rho^i_{op}$ of all the basic operations



Extra power consumption (in Watts)

500

t^{*i*}_{*op*}: RAM Logging (left) and HDD Logging (right)

Calibrate $t_{logging}^{i}$ for each node *i*: mean logging time and deviation



・ロト ・雪 ・ ・ ヨ ・ ・ ヨ ・ うらる

several cores are HDD logging at the same time, t_{op}^{i} is higher

simultaneous accesses on HDD create I/O contentions.

17/24

• different from RAM logging: no contention.

Validation of the estimation framework

4 HPC applications: CM1 and 3 NAS in Class D (SP, BT, and EP) Running over 144 processes (i.e. 12 nodes with 12 cores per node) Measure execution with and without the high-level operations Compute the average value over 30 measurements. Compare estimations to the energy measurements Checkpoint interval = 120 seconds.

18/24

イロト (得) (ヨ) (ヨ) (ヨ) ()

• BLCR system level checkpointing available in MPICH2

Validation of the estimation framework: Accuracy



Relative differences between the estimated and the measured energy consumption.

The relative differences are low => Energy estimations are accurate $> a^{-19/24}$

Validation of the estimation framework: Energy Estimations



20/24

・ 同 ト ・ ヨ ト ・ ヨ ト

э

Not the same energy from one application to another.

How to select the less energy consuming FT protocol ?



The less energy consuming protocol depends on the application Trade-off between volume of data to log and coordination cost.

21/24

Conclusion

A framework to estimate the energy consumption of FT protocols

22/24

3 families of FT protocols: coordinated, uncoordinated and hierarchical

ECOFIT relies on:

- an energy calibration of the execution platform
- a user description of the execution settings

Can be used in any energy monitored supercomputer

Conclusions and Future works

The energy estimations provided by ECOFIT are accurate. Relative difference between estimations and energy measurements:

- equal to 4.9 % in average
- do not exceed 7.6 %.

Select the best FT protocol without pre-executing the application.

23/24

・ロト ・得ト ・ヨト ・ヨト … ヨ

Will propose energy efficient improvements for FT protocols.

Completely include the estimations of recovery in ECOFIT.



Thank you for your attention.



mehdi.diouri@ens-lyon.fr

24/24

< ロ > < 同 > < 回 > < 回 > < 回 > <

э