Parallel Computing 000 (2017) 1–13

Contents lists available at ScienceDirect

Parallel Computing

journal homepage: www.elsevier.com/locate/parco

Using virtualization to quantify power conservation via near-threshold voltage reduction for inherently resilient applications☆

Li Tan*, Nathan DeBardeleben, Qiang Guan, Sean Blanchard, Michael Lang

Ultrascale Systems Research Center, Los Alamos National Laboratory, USA

ARTICLE INFO

Article history:
Received 25 October 2016
Revised 22 May 2017
Accepted 26 July 2017
Available online xxx

MSC:
68M14
68M20
68N25
68W10
68W15

Keywords:
Inherently resilient applications
Power savings
Fault injection
Failure rates
Soft errors
Virtualization
Near-threshold voltage reduction

ABSTRACT

Power efficiency nowadays is a mainstream pressing issue in High Performance Computing (HPC), due to limited power supply capability of current and projected supercomputers. As a promising solution, leveraging inherent application resilience to relax power requirements of HPC runs can effectively save power with minor/acceptable loss of output quality. However, the challenges of this approach lie in: (a) how to reduce power usage of HPC runs online within allowable maximum extent such that quality metrics of applications can be satisfied, and (b) how to identify potential intrinsic nature of fault tolerance in general for an application. Existing efforts to date fail to address both challenges systematically and efficiently. In this work, based on virtualization and near-threshold voltage reduction techniques, we propose an empirical framework named `V-Power` to save the most power for inherently resilient applications. As an integrated empirical system, our approach effectively addresses the two above challenges using quantitative and fine-grained application inherent resilience analysis and frequency-independent near-threshold voltage reduction. Experimental results for a wide spectrum of scientific applications running on a 40-core power-aware server demonstrate that `V-Power` is capable of saving power up to 12.3%, resulting in a failure rate with acceptable program outputs.

1. Introduction

Considering ever-growing power bills and limited power supply capability of electrical facilities, power efficiency has already become a first-class citizen in scalable and cost-efficient High Performance Computing (HPC) systems. Numerous software/hardware solutions have been proposed to achieve power savings at different levels of system abstraction, with little performance degradation of HPC runs. Provided a growing number of specialized-purpose scientific applications in different domains nowadays, a promising solution is to leverage specifically *inherent resilience* of an application for operating in the low-power mode, while the resulting errors are *masked* or *tolerated* by the application. Note that the inherent resilience can be specific to certain error types, or be specific to certain portions of the application. Generally, HPC runs with masked

☆ This work was performed at the Ultrascale Systems Research Center (USRC) at Los Alamos National Laboratory (LANL), assigned the LANL identifier LA-UR-16-23584.

* Corresponding author.
E-mail address: darkwhite29@gmail.com (L. Tan).

http://dx.doi.org/10.1016/j.parco.2017.07.009
0167-8191/

errors result in successful completion, which highly depends on the application characteristics. For instance, the faulty values caused by the errors are not reused in the rest of a run [1]. On the other hand, errors in some applications (e.g., from the domains of image/video processing, recognition, and data mining) can be tolerated to some extent. Incorrect results in a run of such applications within a given range are acceptable [2] (e.g., minor image errors can be unperceivable by users). These applications can benefit from being operated in the low-power mode for power savings, with little loss of output quality.

Strategies at various layers of abstraction can be applied to leverage the inherent resilience of applications, including using inexact hardware [3,4], voltage scaling [5,6], load value approximation [7,8], and task skipping [9,10]. However, little work has been done to investigate the quantitative correlation between resilience and power efficiency of HPC runs, e.g., the integrated modeling between resilience factors such as failure rates, with power saving factors such as supply voltage. In this work, we propose to achieve power efficiency for scientific applications with inherent resilience to different types of soft errors including crashes, hangs, and silent data corruption, using near-threshold voltage reduction techniques [11,12]. There exist several efforts on employing voltage scaling techniques for power savings with minor loss of output quality at different levels of abstraction [5] and at GPU level [6], but neither of them discuss the quantitative modeling of resilience and power savings, nor experimentally evaluate the trade-offs between output quality loss and power efficiency.

CMOS-based hardware components such as CPU, GPU, and memory are generally the dominant power consumers in a computing system. Lowering supply voltage of them can effectively reduce power costs of the system, given that power consumption of these components is proportional to the product of operating frequency and supply voltage squared [13], which means voltage reduction has a greater impact on power savings than frequency reduction. Existing efforts extensively demonstrate that Dynamic Voltage and Frequency Scaling (DVFS) techniques can effectively save power/energy when the peak performance of hardware components (CPU [14,15], GPU [16,17], and memory [18,19]) is not necessary. Nevertheless, a critical limitation of existing DVFS work lies in the fact that the employed DVFS techniques are essentially *frequency-directed*: Voltage and frequency are scaled down together in the presence of hardware idle time [20], which fails to fully exploit potential power saving opportunities during HPC runs. As a remedy solution, *undervolting* [12,20,21] has been employed to further reduce power costs, which reduces voltage independent of frequency scaling. Using undervolting, hardware components are supplied with a voltage that is lower than the one paired with a given frequency used in DVFS.

As a generic approach applicable during any stages of HPC runs, undervolting is capable of saving extra power in addition to using DVFS during hardware idle time. Nevertheless, the trade-off of using undervolting is increased failure rates [11] of the components with lowered voltage. For inherently resilient applications, the growing number of errors can be tolerated to some extent, and thereby incur negligible impacts on performance and output quality. Therefore, significant power efficiency can be achieved with minor (or acceptable, depending on application characteristics) loss of resilience. In this work, targeting inherently resilient applications, we leverage the resilient nature of specific types of applications and frequency-independent near-threshold voltage reduction, proposing a framework of Virtualization-based `Power`-efficiency (`V-Power` for short) to save the most power with minor/acceptable loss of output accuracy. We aim to explore the fact that, theoretically and experimentally, a wide scope of (and extensively used) scientific applications running on future exascale systems can benefit from power efficiency, if resilience is provided to some extent as the inherent application characteristics. In summary, the contributions of this paper include:

- We leverage inherent resilience of applications for considerable power savings with minor/acceptable loss of program output quality, using frequency-independent near-threshold voltage reduction;
- We develop an integrated virtualization-based framework named `V-Power` to analyze potential resilience of applications, and achieve power efficiency for a runtime-dominant kernel/function of the application, if it is validated resilient by nature;
- Our approach is experimentally evaluated on a power-aware platform for a wide scope of scientific applications to save up to 12.3% in power, with an error rate for acceptable program outputs.

The remainder of the paper is organized as follows: Background knowledge is introduced in Section 2. Details of the proposed `V-Power` framework are presented in Section 3. Experimental results are provided in Section 4. Section 5 discusses related work and Section 6 concludes.

2. Background

2.1. Inherently resilient applications

Regardless of hardware relaxation techniques [22,23], some applications are inherently resilient to errors, i.e., for provided inputs, when errors in runs occur, they produce *acceptable* outputs with (partially) incorrect or inaccurate results. For instance, pixel errors in an image or a video can exist to some extent without the perception from users (or with user-acceptable quality degradation). For some approximate computing applications, a range of output values are considered acceptable, since there is no unique golden output for a golden run. Due to precision requirements of some iterative scientific applications, computation errors occurring in one iteration can get cancelled by truncated numbers during the computation in later iterations [24]. This property of self-tolerance/self-healing to faults is referred to as *inherent application resilience*, arising in some categories of scientific applications. Specifically, applications with such a property span from domains of image/video processing, recognition, and data mining. Notably, the acceptance of program outputs is defined

Please cite this article as: L. Tan et al., Using virtualization to quantify power conservation via near-threshold voltage reduction for inherently resilient applications, Parallel Computing (2017), http://dx.doi.org/10.1016/j.parco.2017.07.009