Adapting Fault Resilience Granularity to Overcome Varying Failure rates in CPS

Laura Rozo, Jose Monsalve, and Chengmo Yang
University of Delaware, Newark, DE
Email: {lrozo, josem, chengmo}@udel.edu

Keywords. Cyber-Physical Systems, Fault Resilience, Adaptation, Energy Efficiency.

I. Background

Cyber Physical Systems (CPS) bridge the physical and the cyber world through using computational components to coordinate, monitor and control the behavior of physical processes in the real world [1]. They are widely employed in various application domains that have great technical, economic, and societal impacts, including medical care and health (e.g. pacemakers, infusion pumps), transportation and mobility (e.g. vehicle based safety systems, ABS, traction and stability control), manufacturing (e.g. computer controlled machine tools and equipment), among others. Given these diverse application domains, an increasing demand of more complex functionality, higher performance, and lower energy consumption is expected. While increasing the number of processing nodes and shrinking the device sizes are desirable, they both have adverse impacts on reliability. More processing nodes imply more resources that may fail, while smaller devices imply higher temperatures and higher vulnerability to radiations.

Since most CPS applications are mission critical, it is mandatory for these systems to deliver higher reliability and guarantee correct functionality even in the presence of failures due to battery depletion, physical damage, or environmental interference. Implementation of effective fault tolerance techniques plays an important role. What makes this problem more difficult than fault tolerance of general purpose systems are the strict requirements that Cyber Physical Systems usually have, including severe energy constraints, distributed control, and complex network interactions between controllers.

Current fault tolerant techniques for CPS are either specialized to particular applications [2]–[4] and cannot be employed in areas that do not necessarily share exactly the same characteristics, or targeting specific faults, such as the ones caused by security attacks [5], and cannot cover faults due to malfunction processors or environmental effects.

In comparison, we propose an adaptive fault tolerant scheme that is independent of the type of applications and the type of faults. Our scheme efficiently overcome the reliability challenge through adapting fault detection and recovery granularity based on the available resources and fault rate conditions of the system. This way, the performance and power overhead imposed by fault tolerance can be minimized.

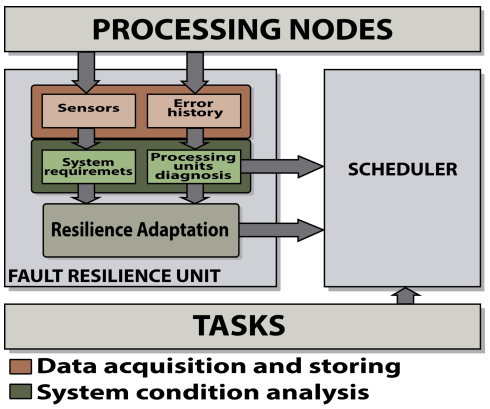

Fig. 1. Overview of the proposed Fault Tolerance Framework

II. Proposed Project

The main idea of the proposed fault tolerant scheme is to adapt the fault detection and recovery process according to the current conditions of the system. These conditions include power consumption and fault rates, which will be monitored by sensors and specialized units. Figure II illustrates the functional overview of the proposed work. *Fault Resilience Unit* is the component that coordinates all the fault tolerant functions proposed in the project. Its functionality will be divided in three parts: data acquisition and storing, system condition analysis, and resilience adaptation. *Data acquisition and storing unit* recollects and stores important information in the system, including the temperature and power information of each processing node, as well as the error history of all of them. *System condition analysis unit* provides three main functions. It verifies that system requirements are being satisfied (i.e. power constraints), identifies the unhealthy nodes and informs the scheduler to avoid binding tasks to them, and computes current fault rate of the system. Finally, the *resilience adaptation unit* is responsible for making decisions regarding which fault detection and recovery schemes should be used, based on the system condition information obtained from the previous unit.

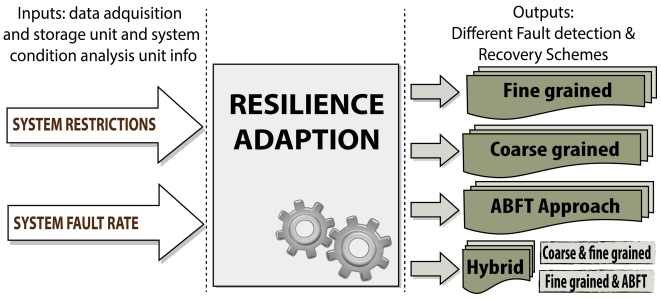

Fig. 2. Functionality of the Resilience Adaption Unit

Various fault detection and recovery schemes can be employed by the resilience adaptation unit to maximize overall efficiency. Figure 2 shows a detailed functional view of this unit. As can be seen, four different fault detection schemes can be used, including *fine-* and *coarse-grained* schemes, *algorithm-based* schemes and hybrid schemes.

Both *fine-* and *coarse-grained* schemes exploit hardware redundancy wherein each task/instruction is redundantly executed on different nodes to detect faults, if any, and recover the affected computation. Yet they differ in the fault detection and checkpointing granularity. In the *fine grained* scheme, fault detection and checkpointing are performed very frequently (e.g. each time that an instruction is executed). When the fault rate is high, this approach allows faults to be detected right away, thus preventing them from affecting subsequent computations, however, at the cost of being more time, energy and resource consuming. In comparison, *coarse grained* schemes perform fault detection and checkpointing less frequently (e.g. each time that a whole task is executed), thus imposing less overhead in performance, energy and resources. This approach is more suitable when the fault rate is low.

The *Algorithm Based Fault Tolerance (ABFT)* scheme [6] is suitable for systems that need to process a large amount of data. Instead of duplicating the data set, this approach allows the use of single data structures and performs fault detection and correction by augmenting those data structures to include additional information, such as parity bits, checksums, etc. This scheme is more efficient in terms of time, energy and resources, compared to fine and coarse grained schemes. However, it is only applicable to computations with high regularity.

Finally, the resilience adaptation unit can also employ different schemes for different tasks or at different stages of the computation. The most appropriate approach will be selected based on the priorities specified by the designer, as well as the current system conditions monitored by the sensors. Designer can define thresholds for fault rate and power consumption and assign different priorities. For example, if both fault rate and power consumption exceed the threshold, priorizing reliability over energy would allow the system to adapt to more effective resilience schemes, and be able to sacrifice performance and power constraints to some extend, avoiding potential system crash due to high fault rates. On the other hand, priorizing energy constraints over reliability, in the same scenario, would lead to the use of a more coarse grained resilient scheme that is less energy, time, and resource consuming, yet sacrificing the system's ability to quickly detect faults and recover from them.

III. Impact in Cyber-phisical Systems

The proposed scheme is expected to have transformative impact on fault resilience in cyber physical systems. Previously fault resilient systems were developed to work in a static environment wherein fault rate is low and variations in the fault rate are negligible. However, in Cyber-Physical Systems, the computational nodes will constantly interact with the physical world, causing high fault rate and diverse fault behavior due to various reasons such as battery depletion, physical damage, and environmental interference. While creating an aggressive fault resilience scheme that covers the worst case scenario across system lifetime is theoretically possible, this strategy is unpractical since it will engender significant overhead that violates other CPS design requirements such as energy and workload constraints.

With the fault resilience solutions presented in this document, the overhead can be largely reduced while the system is allowed to tolerate a diverse range of fault rates. In summary, our work can help future CPS, especially those mission critical applications, to survive the increasing amounts and diverse types of hardware failures. This in turn, allows CPS to operate in harsh environments. It also relaxes the requirement of 100% correctness for devices and interconnects, thus dramatically reducing the costs of manufacturing, verification, and testing.

References

[1] W. Wolf, "Cyber-physical Systems," *Computer*, vol. 42, no. 3, pp. 88–89, 2009.

[2] K. Sampigethaya and R. Poovendran, "Cyber-physical system framework for future aircraft and air traffic control," in *Aerospace Conference, 2012 IEEE*, March 2012, pp. 1–9.

[3] C. Krishna and I. Koren, "Adaptive fault-tolerance for cyber-physical systems," in *Computing, Networking and Communications (ICNC), 2013 International Conference on*, Jan 2013, pp. 310–314.

[4] C. Singh and A. Sprintson, "Reliability assurance of cyber-physical power systems," in *Power and Energy Society General Meeting, 2010 IEEE*, July 2010, pp. 1–6.

[5] G. Sabaliauskaite and A. Mathur, "Intelligent checkers to improve attack detection in cyber physical systems," in *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2013 International Conference on*, Oct 2013, pp. 27–30.

[6] I. Koren and C. M. Krishna, *Fault-Tolerant Systems*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2007.