

DOME: Delaying and Overcoming Microprocessor Errors

Negar Miralae | Jyothish Soman | Timothy Jones | Alan Mycroft | Computer Laboratory, University of Cambridge

Project Overview

The project explores the mechanisms required, both in hardware and software, to delay the occurrence of errors and, once the error has occurred, to continue working while tolerating those errors.

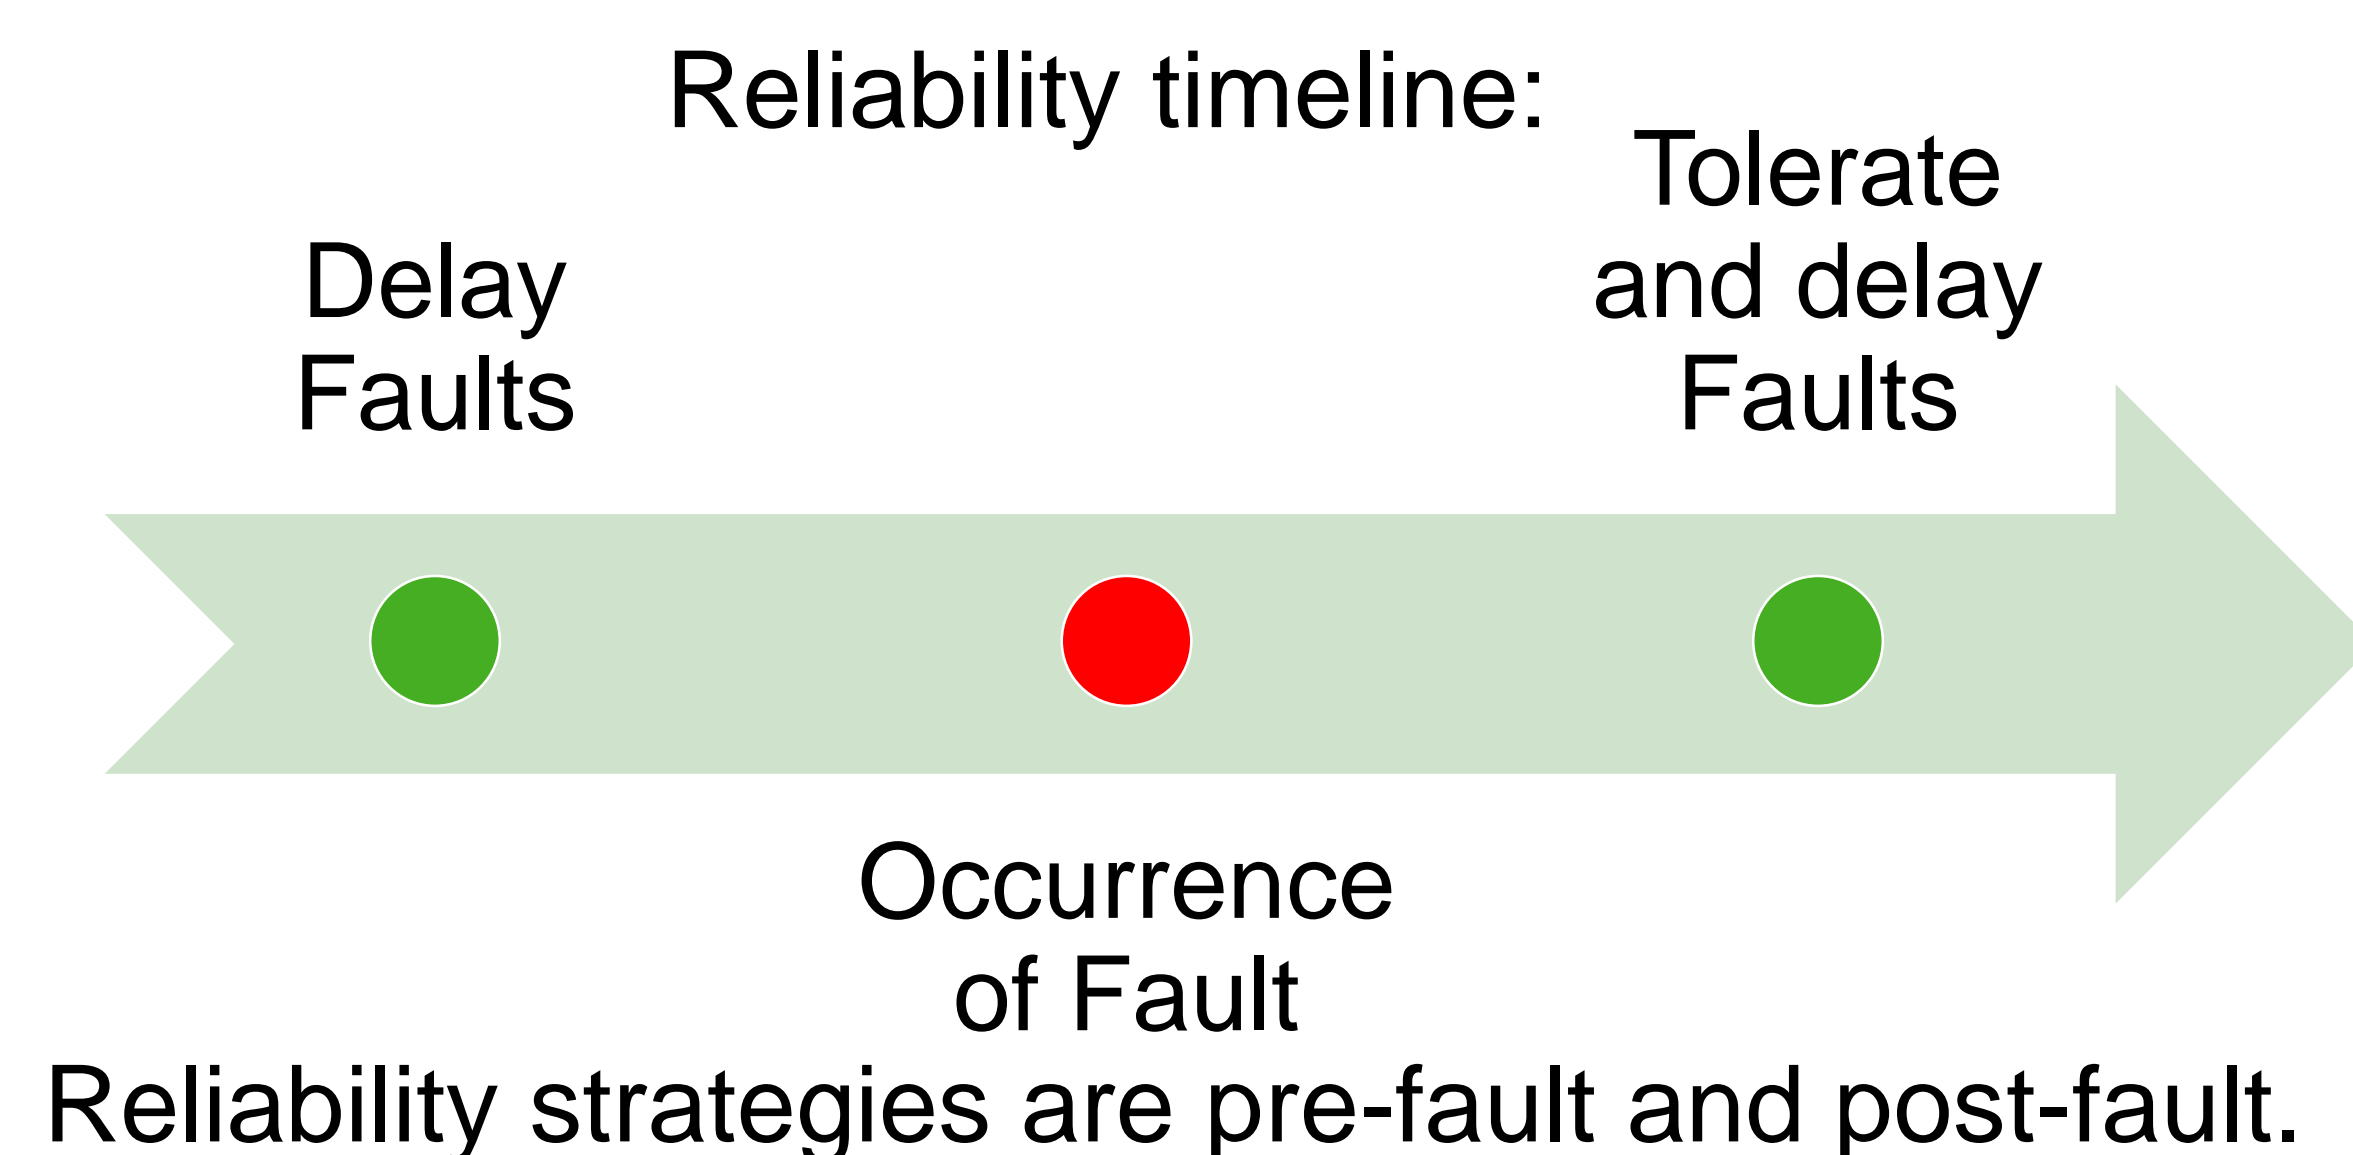
The project started in September of 2012 and will finish in April of 2016. This project is a collaboration between the University of Manchester and the University of Cambridge.

Background

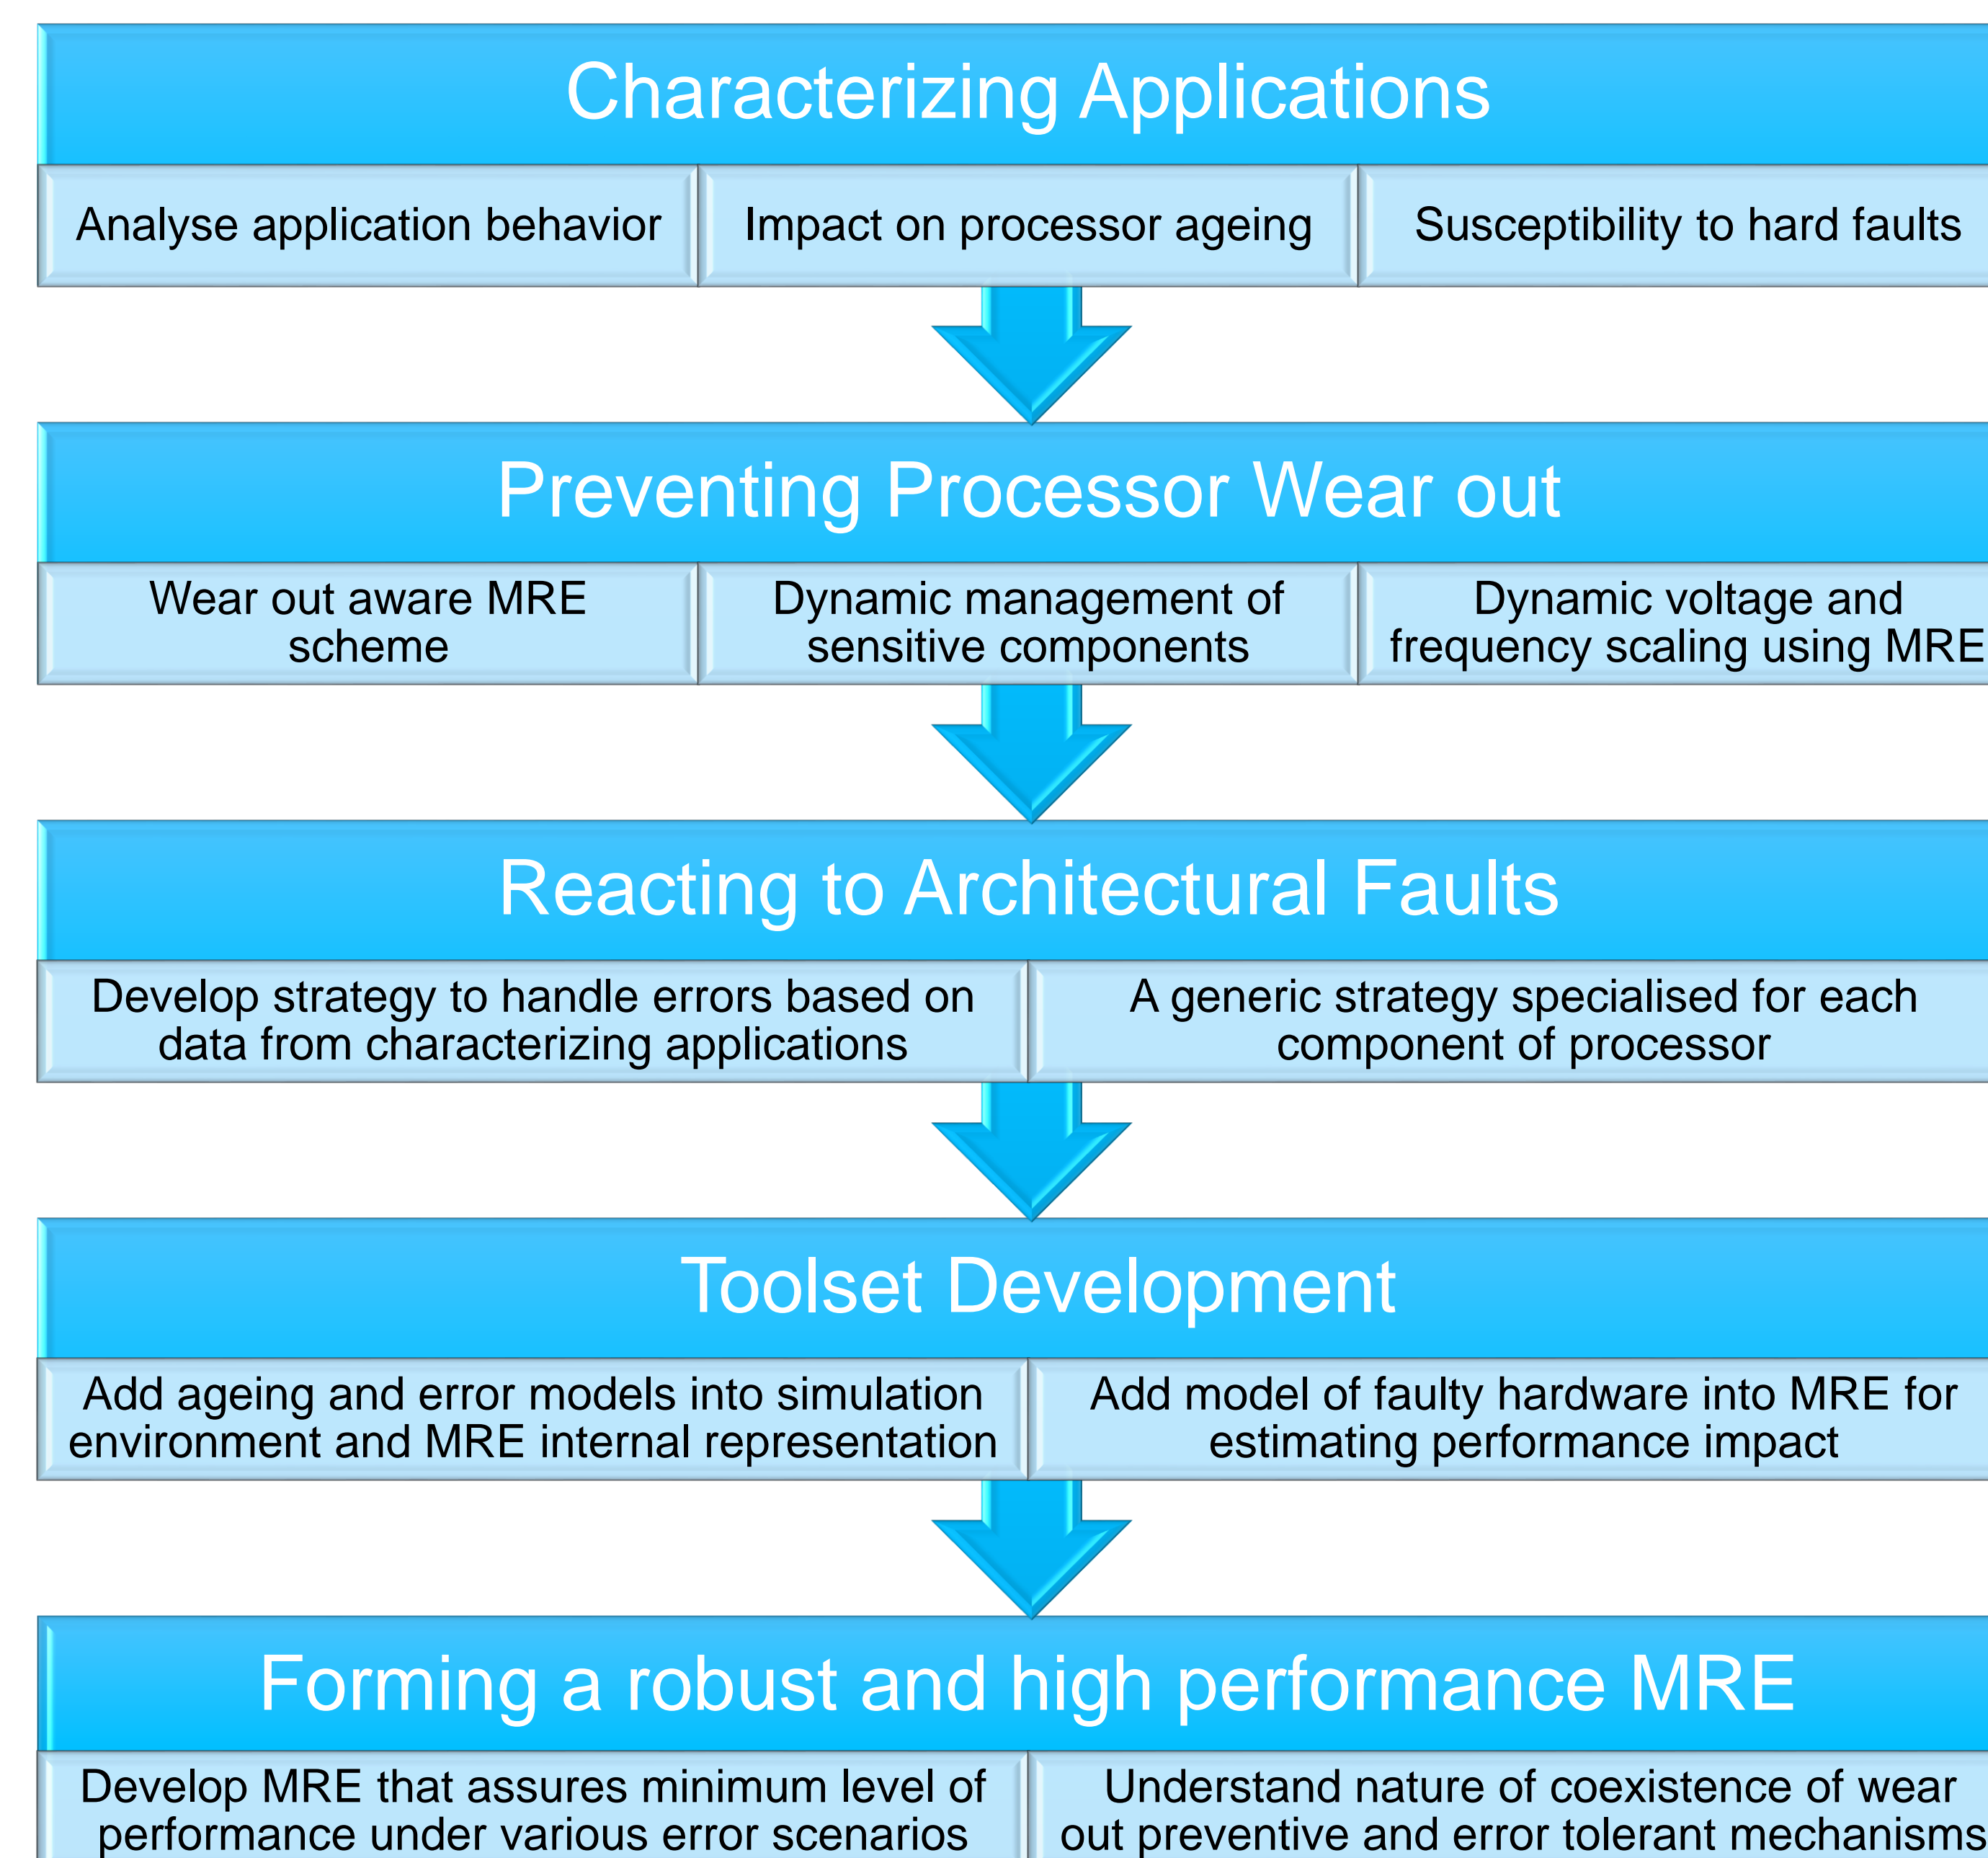
Power inefficiency of current transistors has led to an increased focus on power management, making traditional error handling schemas such as redundant hardware increasingly infeasible. Additionally transistors are taking less time to wear out, thus becoming more prone to errors. The trend of reducing transistor sizes and increasing their number on a single chip, makes reliability challenges more critical than ever. Hence alternate designs of hardware and software are needed.

We are taking a novel approach to these challenges by leveraging managed runtime environments (MRE). MREs can be made aware of the wear out and faulty behaviour of the processor, assisting runtime monitoring and altering of applications to increase chip lifetime.

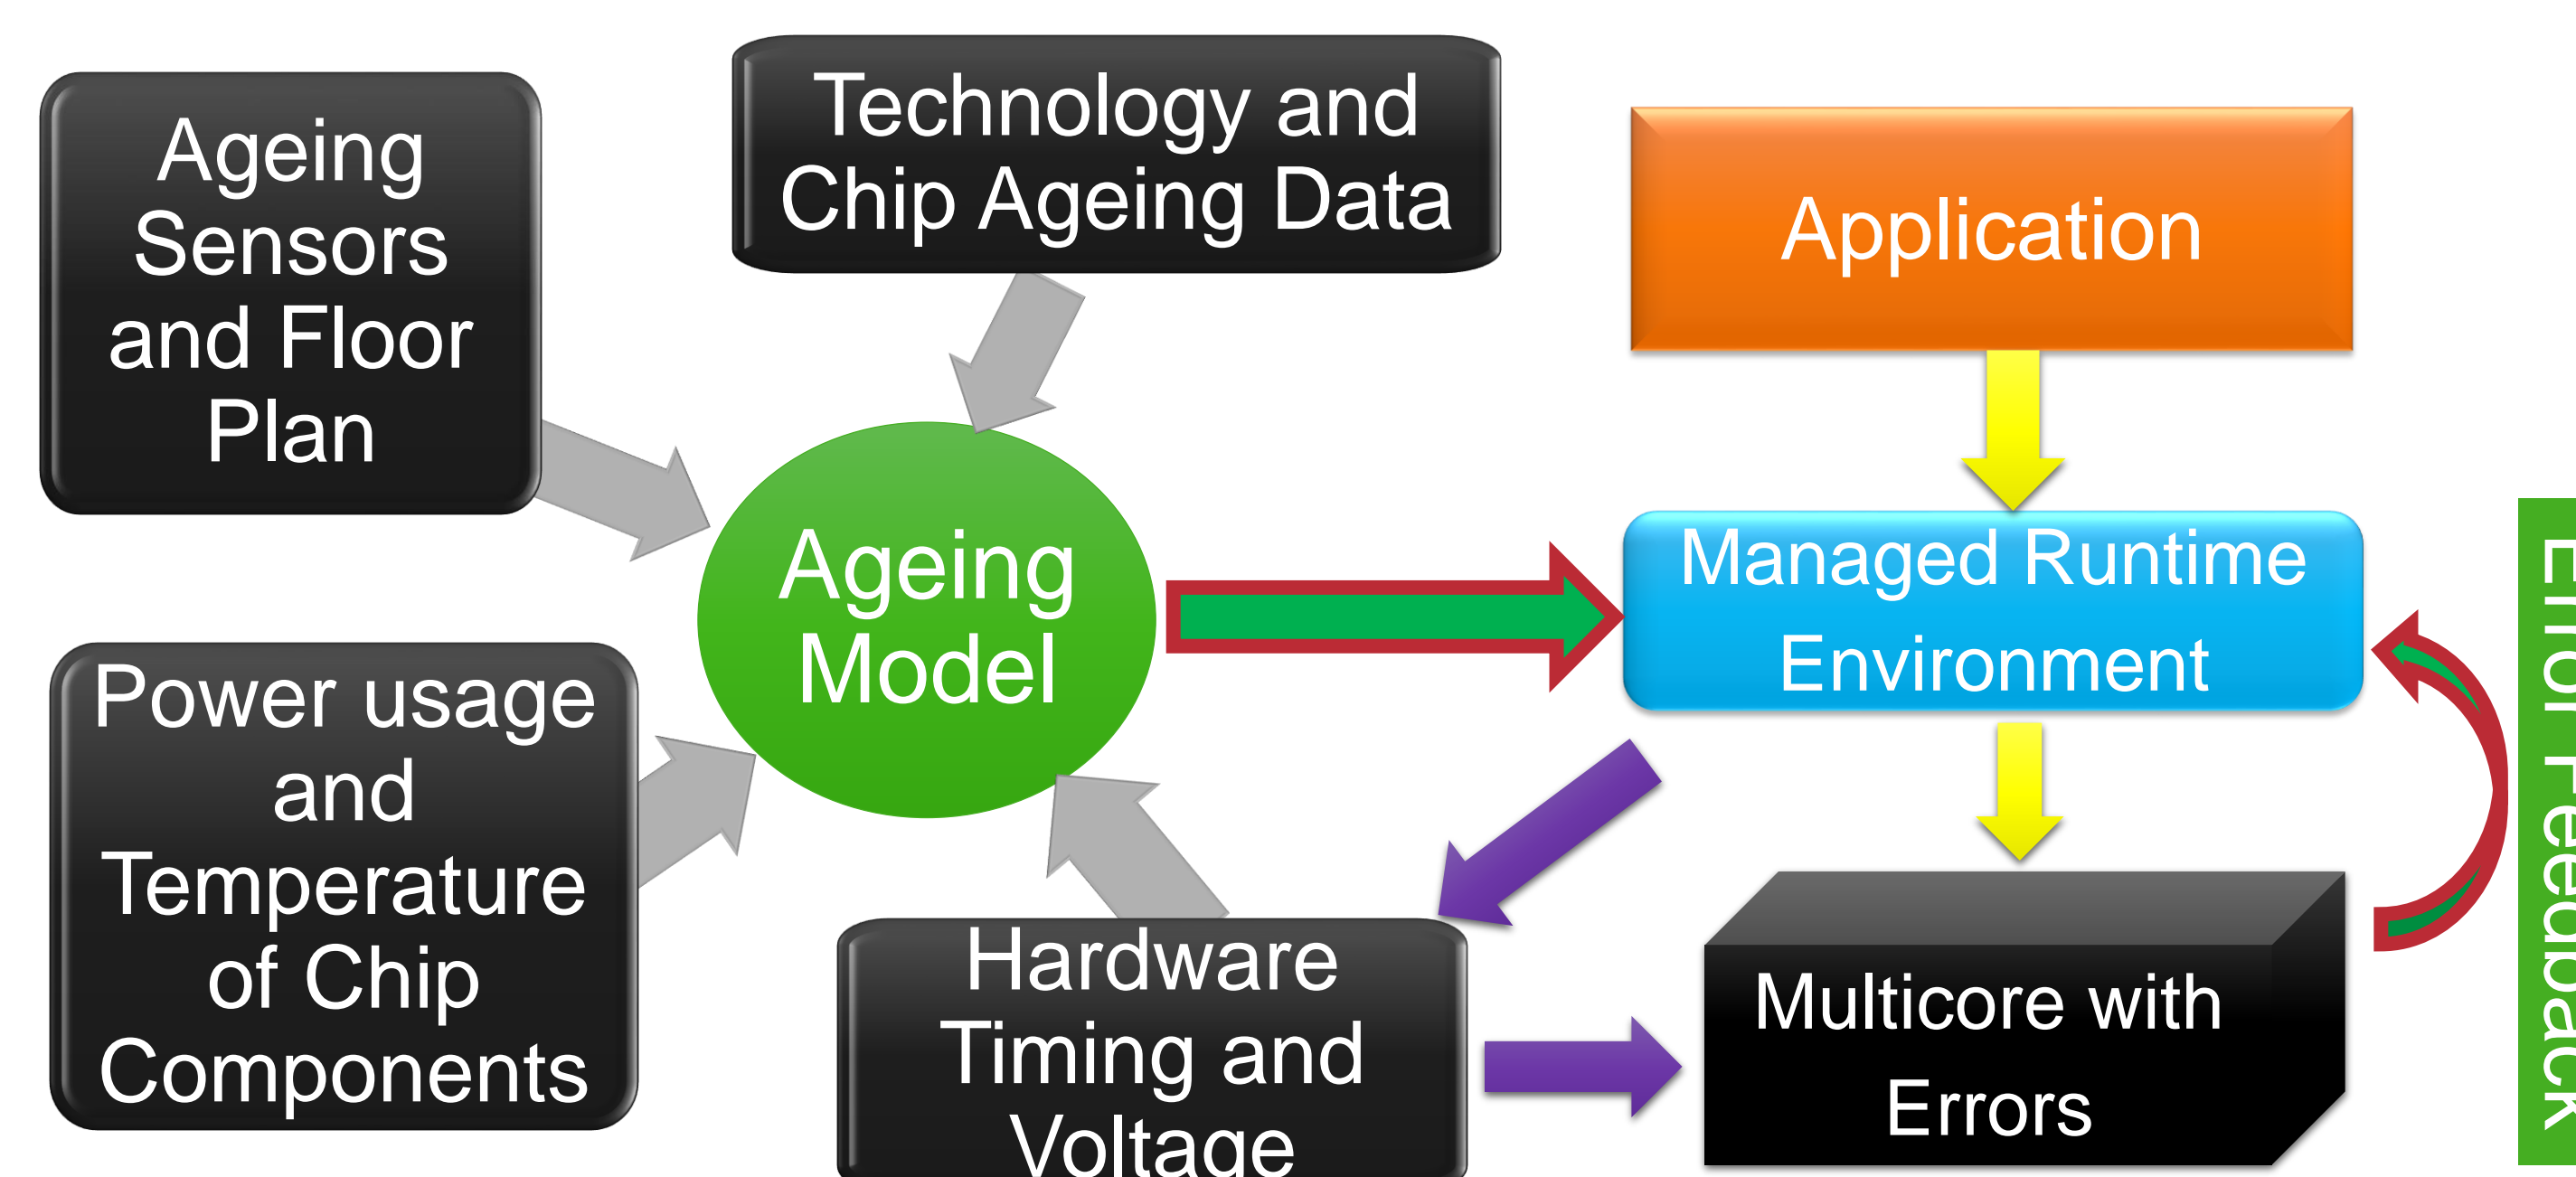
Managing Reliability



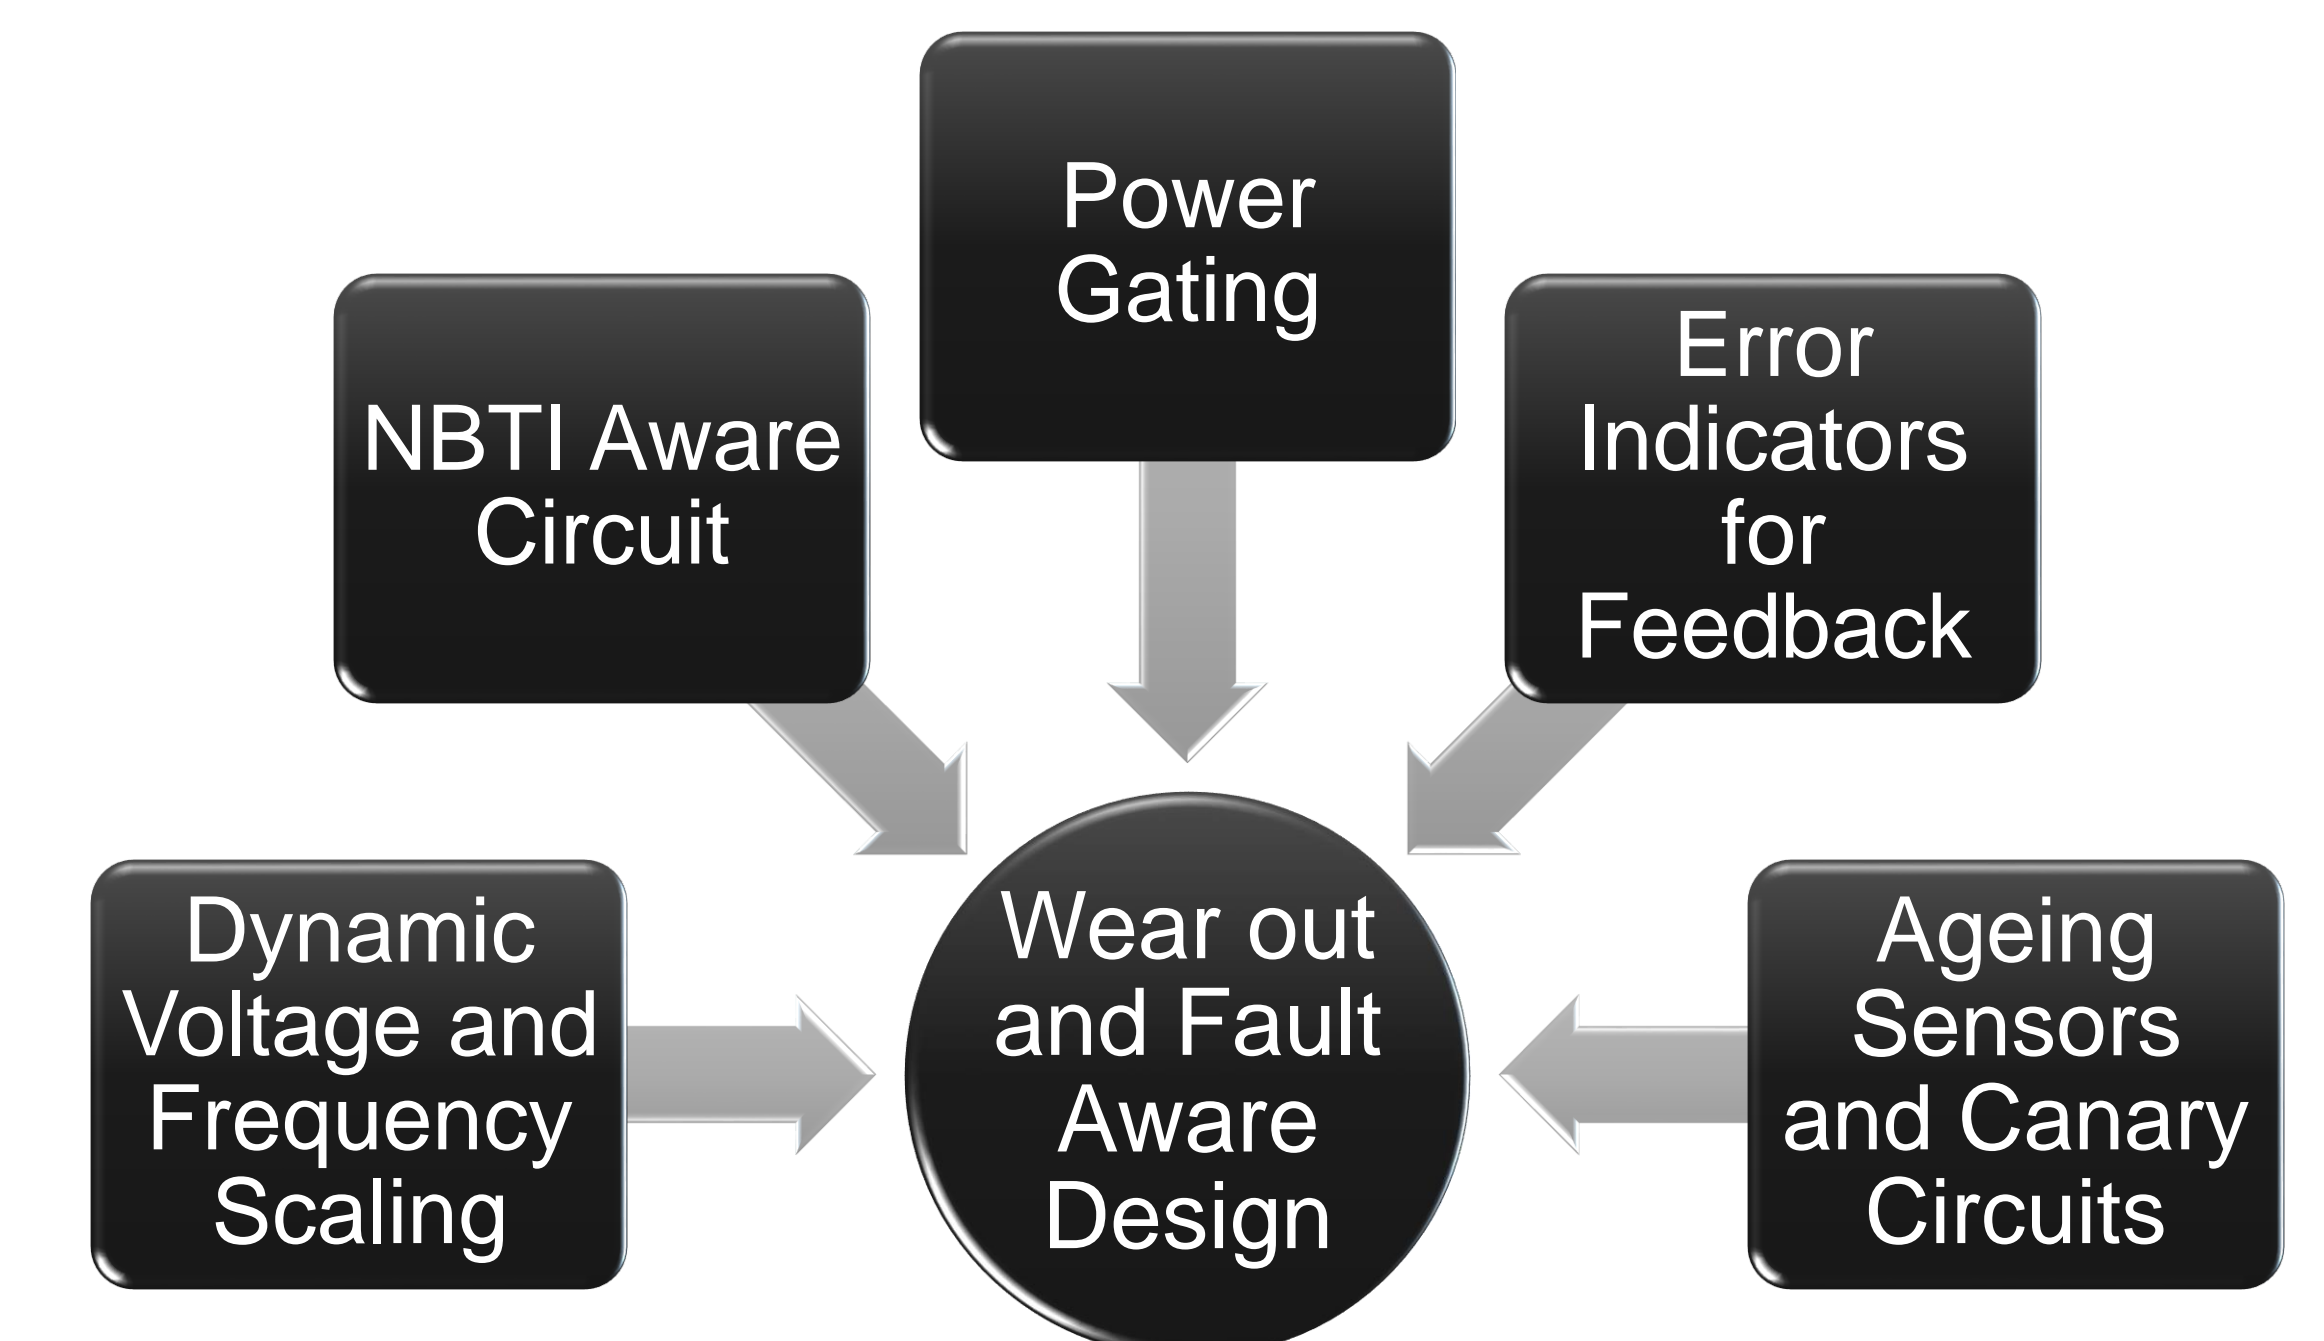
Timeline of Project



Software Architecture

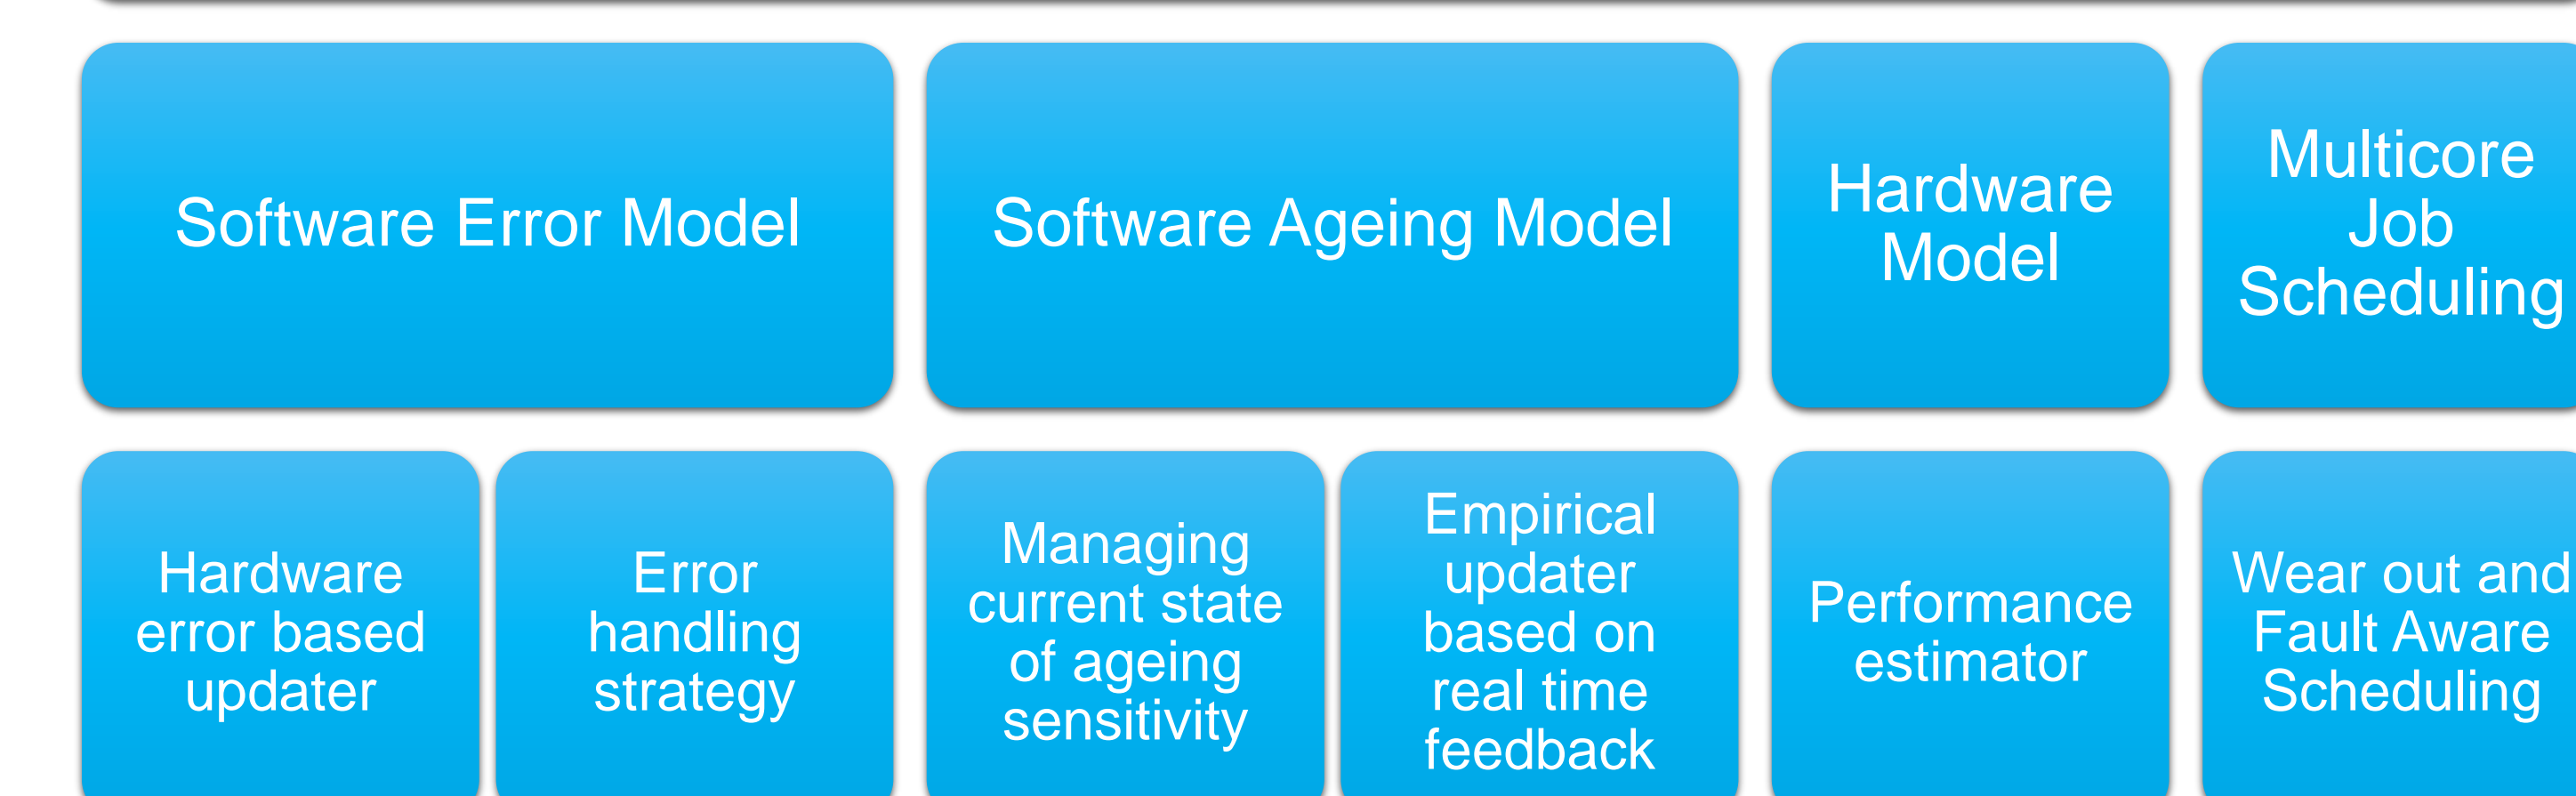


Hardware Design Directions



MRE Architecture Overview

Managed Runtime Environment



Generalised Aims of the Project

- Our aim is to create a general wear out and fault-aware framework that is suitable for most processor architectures.
- The generality extends also to the applications ran on the system. The project though also focuses on finding general trends in behavior of the system across various applications which can assist in creating a practical design.