




## Studentship Project

### Non-invasive 'safety agents' for embedded processors

Michael J. Pont and Kam L. Chan, University of Leicester

As embedded designs take on an increasing role in system safety, it becomes important that we make our embedded processors operate as predictably as possible and accept that no embedded system will ever be 100% reliable and provide cost-effective ways of ameliorating the impact of processor and programming errors when they do occur. We have demonstrated in previous studies that using "time triggered co-operative"(TTC) architectures can lead to extremely predictable system behaviour at low cost (e.g. Phatrapornnant and Pont, 2006). We have also demonstrated a novel "time triggered cooperative" (TTC) processor which can help to ensure that TTC architectures are applied correctly. The combination of high predictability and low cost produced by such architectures is a key consideration in – for example - the automotive sector, where we are seeing major investments in moves towards "drive by wire" vehicles.

We would not seek to claim that TTC processors are perfect: for example, even when programmed correctly, they are (like all embedded processors) susceptible to the impact of electromagnetic interference (e.g. see Ong et al., 2001). Our aim in this project is to further improve the reliability of systems employing devices such as a TTC processor. Specifically, we are exploring techniques which will allow us to monitor the activity of a wide range of embedded processors and to reset the processor (or take other appropriate action) in the event that errors are detected. In this study, our monitoring unit (a "safety agent") is based on a second, very simple, processor node. We accept at the outset that this monitoring process presents a significant challenge. Put simply, we would ideally like to separate the safety agent from the main processor, in order to reduce the possibility that (for example) electromagnetic interference will affect both the main processor and the agent. However, without contact between the two devices, the monitoring process itself becomes a practical impossibility. This project is exploring an approach to this problem that involves having the safety agent examine the activity of the main processor by monitoring fluctuations in the CPU power consumption. Through this means, we aim to achieve a very high degree of separation between the main processor and safety agent and still provide an effective monitoring operation.



**Figure 1: Monitoring in progress.**

The project is progressing very well. We have been able to detect a number of errors through measurements of system power consumption (see Figure 1). A first patent has been filed and – with an industrial partner (TTE Systems Ltd) - we are discussing development of a commercial product based on this technology.