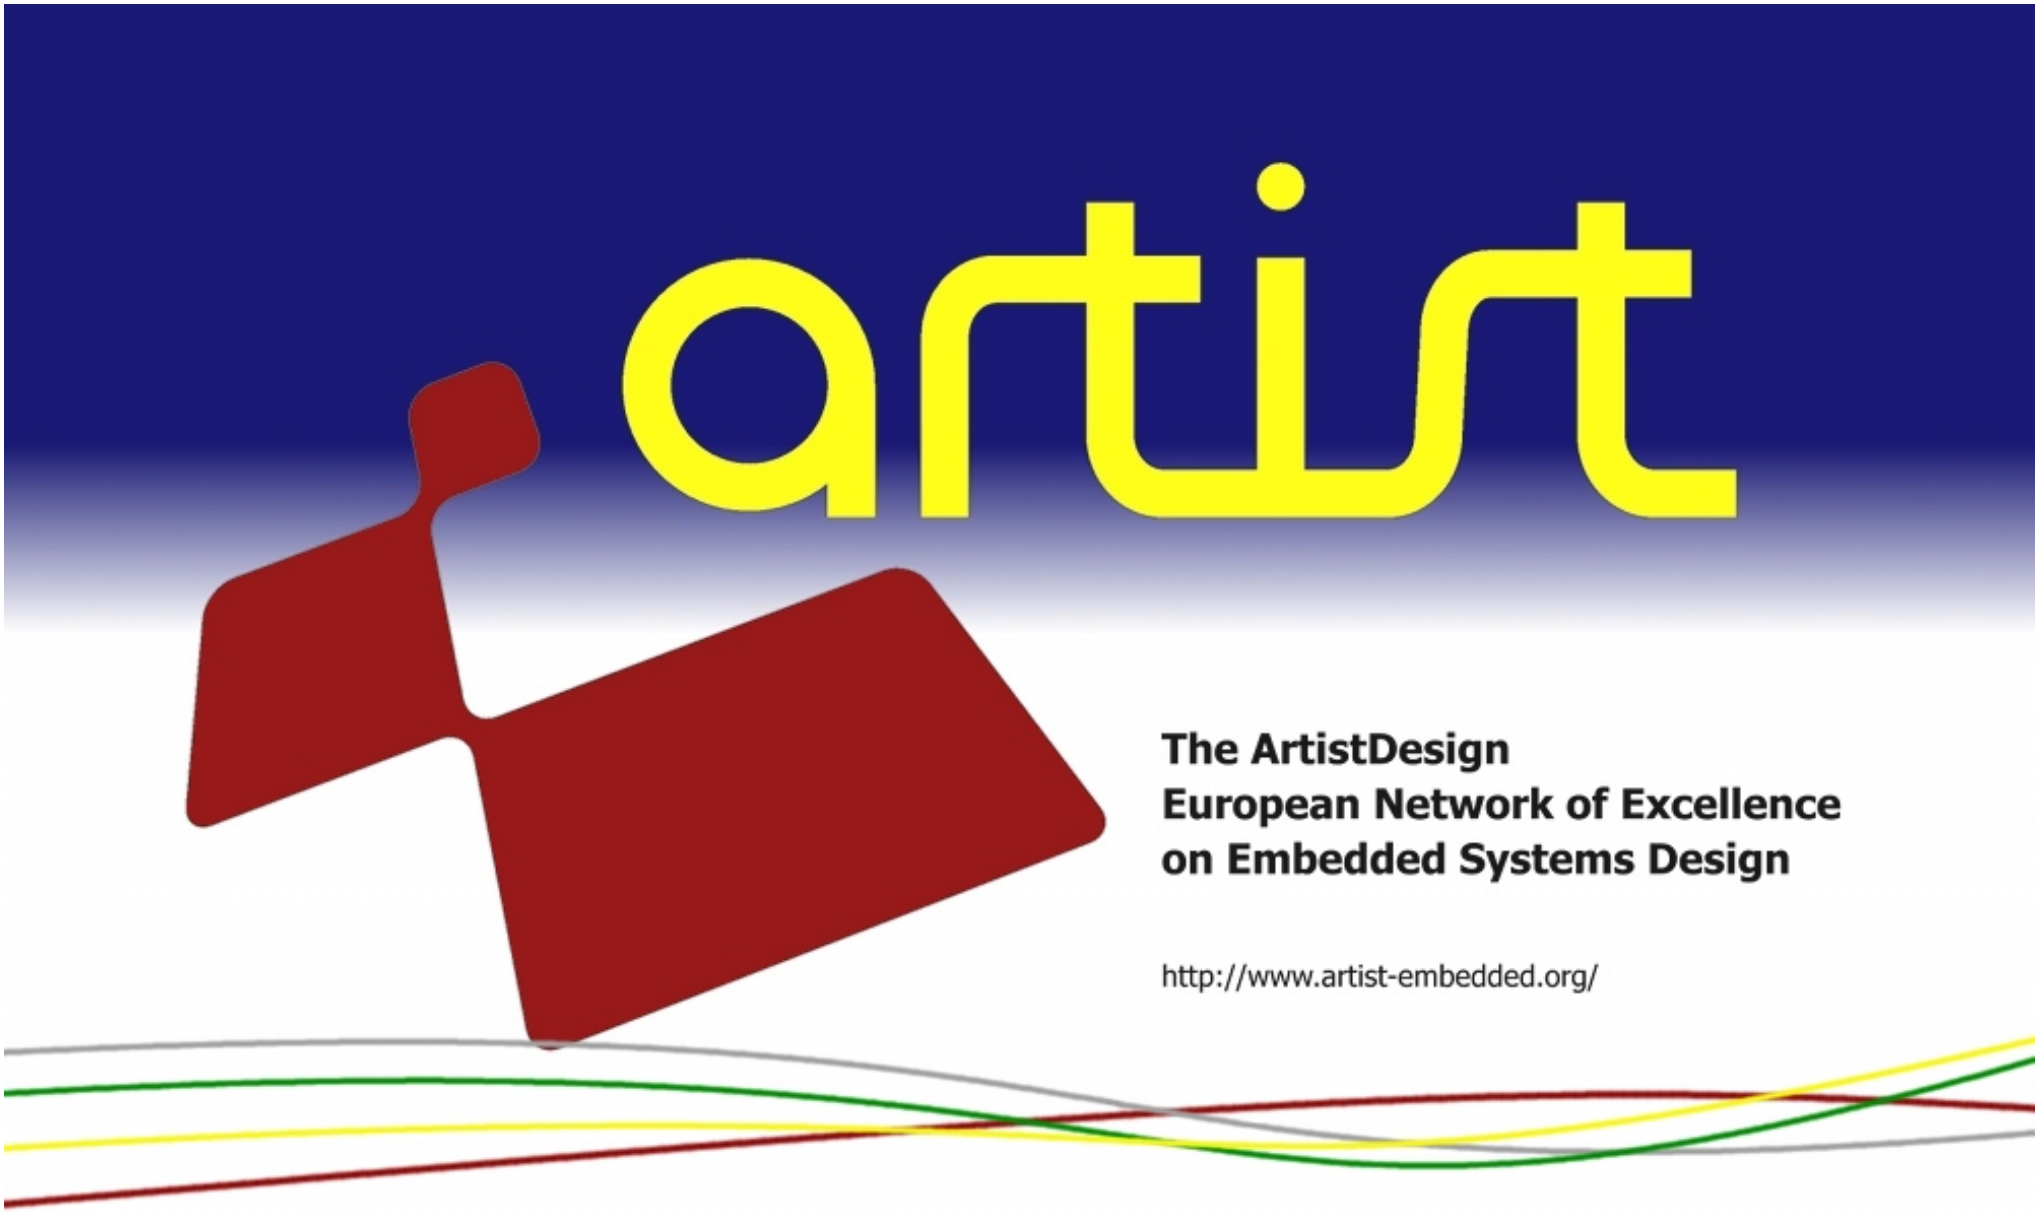

Showcase of the Main Results

DATE Conference, March 15th, 2012

SEVENTH FRAMEWORK PROGRAMME

ArtistDesign European NoE: Showcase of the Main Results

DATE Conference, March 15th, 2012

Achievements and Perspectives :

Scheduling and Resource Management

Alan Burns

University of York, UK

Objectives

- **Provide Policies**
  - For effective resource usage
- **Provide Analysis**
  - For predicting system behaviour
  - Simulation, scheduling analysis, measurement, model checking
- **Provide Models**
  - For composing systems
  - Time triggered and event-triggered work flow
  - For static and dynamic usage patterns

Challenges

- To move from single processor platforms to multiprocessor, multi-core, FPGA, etc.
- To integrate various resources and abstract views of the overall system
  - Integrate policies
  - Integrate analysis
  - Integrate models
- Static and Dynamic, peer-to-peer and hierarchical

Outputs

- In four years, ArtistDesign partners have produced
  - 92 “Technical Achievements”
  - Over 400 refereed papers
    - Including a major review of multiprocessor scheduling published in ACM Computer Surveys (2011)

Results

- Significant work still on single processor systems, for example
  - Efficient analysis for EDF
  - Energy and power aware scheduling
  - Sensitivity analysis and sustainable analysis
  - Parameters selection for control systems
  - Limited preemptions
  - Optimality results

Results

- Language and other standards work
- Much work on contract-based (virtualisation) means of integrating components
  - Recently extended to mixed criticality systems
- Hierarchical scheduling of various forms
- Distributed Systems
- Multiprocessor scheduling

Multiprocessor Scheduling

- For globally allocations:
  - Better priority assignment (Deadline Monotonic is far from optimal)
  - For EDF and Fixed Pri – schemes that switch to least laxity at some point (eg. EDZL, FPZL)
  - Better scheduling tests for Fixed Pri – though some not compatible with optimal priority assignment
  - No optimal scheme for sporadic task sets (without clairvoyance),
- Overheads – Good News (migration = preemption), Bad News (shared queues etc prohibitive for N > 6)

Multiprocessor Scheduling

- For fully partitioned we still have the 50% bound, but
  - For systems of small tasks, schemes such as first-fit on density work well (largest density first)
- Semi-partitioned approaches are proving to be more useful
  - What is the minimum number of migrations to get optimal performance (if cost of migration and preemption is ignored)
  - What is the best performance we can get from a one-task-per-core migration scheme

Task Splitting

- Most tasks are statically allocated, N-1 are split between processors (for N CPUs)
- One task splitting scheme for EDF scheduling has a task split (C, D, T) so that first part has C1=D1, C1<C
- The second part (C-C1, D-D1, T) then has maximum time to execute on second processor
- Often 100% utilisation is achievable (when overheads are ignored)
  - But overheads are potentially very low
- General performance is very good
- Equivalent scheme for Fixed Pri has been analysed

Open Issues

- As we move to many-cores, the thread/task is no longer the right abstraction on which to partition work-flow
  - So concurrency within tasks must be addressed
- Still no effective resource control protocol for multi-core platforms (for partitioned or global allocation)
- On a multi-core, WCET analysis, task scheduling analysis and NoC analysis must be dealt with holistically
- Mixed criticality then adds to the fun