a whitebox RTOS for multicore embedded systems

GR740 User Day

13th December 2022

Michael Ryan, CTO, O.C.E.Technology Ltd

Embedded System Characteristics (1)

- **Fixed Code Base** *Software not added during system life*
  => can use physical addresses, RTOS doesn't need page tables etc.
- **Code must be robust** *So an RTOS design should…*
  => ensure certain failure modes are impossible
  - e.g. unbounded priority inversion, …

  => allow behaviour policing by the application (white box)
  - performance data recorded and checkable at any time

  => automatically check for problems
  - e.g. stack overrun, missed deadlines…

  => automatically trigger application problem handlers
  => provide calls to deal with problems
  - e.g. kill task, disable task, disable a CPU, …

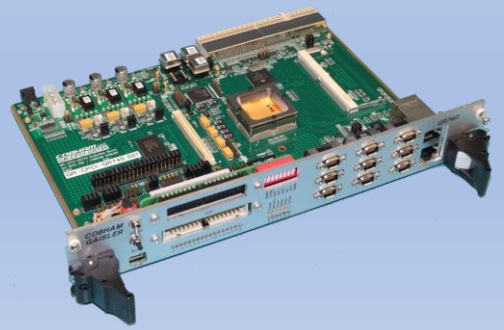

OCEOSmp on GR-CPCI-GR740

Embedded System Characteristics (2)

- **But things can go wrong** *So an RTOS design must provide*
  => **F**ault **A**nticipation, **D**etection, **I**solation, **R**eporting, **R**ecovery
  - make it easy for the application to police the system:
    min time between task start requests
    max execution times, deadline misses
    max pre-emptions, max stack usage
  - make it easy to exploit redundancy:
    run same task on many CPUs, check results agree
  - automatic checks of key components
    memory area sentinels, stack space
  - automatic logging of anomalies
  - triggers for user defined problem handler functions
  - kill tasks, disable tasks, disable CPU cores,
    restart CPU core, …

OCEOSmp on GR-VPX-GR740

OCEOSmp: For Multi-Core Embedded Systems

- **Based on ‘Stack Resource Policy’ (Baker 1991)**
  => single system stack per CPU (not stack per task)
- **Deterministic**
  => behaviour predictable
  => memory statically allocated
  => timing overheads minimized and quantifiable
- **Schedulability analysis**
  => simpler to perform
- **Application task timing recorded for analysis**
  => maximum execution times, maximum times to completion,…
  => missed deadlines trigger application defined action
- **Timed actions independent of scheduling**
  => output at specific time, task start request at specific time

OCEOSmp on GR712

OCEOSmp: Multi-Core

- **Exclude**
  => exclude some cores from OCEOSmp use, e.g. for use by Linux
- **Reserve**
  => reserve cores for higher priority OCEOSmp tasks
- **Symmetric**
  => after start-up, all cores are equal
- **Work Distribution**
  => task execution instances distributed evenly across cores
  => unless task restricted to a particular core
- **Control**
  => take core out of use, put core back in use

OCEOS on GR716

OCEOSmp: RTOS (1)

- **Fixed priority**
  => task priorities fixed based on task importance
- **Pre-emption threshold**
  => pre-emption only by tasks with higher priority than threshold
- **Multiple execution instances**
  => multiple same task ‘jobs’ can be in execution at same time typically using different data
- **Timed actions independent of scheduling**
  => data output at specific time
  => task start request at specific time

OCEOSmp: RTOS (2)

- **Mutexes**
  => unbounded priority inversion cannot occur
  => deadlock warning, cannot occur if single core
- **Read-Write mutexes**
  => allow multiple simultaneous reads of protected area when not being written, prevent writing if being read
- **Counting semaphores**
  => allow wait with timeout
- **Data queues**
  => allow read with timeout

Solar array drive

OCEOSmp: RTOS (3)

- **System time**
  => in microseconds, 64 bit
- **Context switch timing**
  => context switching shared across all cores
  => context switch time minimized
- **Interrupts**
  => interrupt disabled timing is minimized
  => high priority timer interrupt reserved for timed actions
- **Some numbers**
  Up to 255 cores, 255 tasks, 15*255 execution instances (jobs),
  63 mutexes, 63 read-write mutexes, 63 semaphores,
  63 data queues, memory < 20KiB

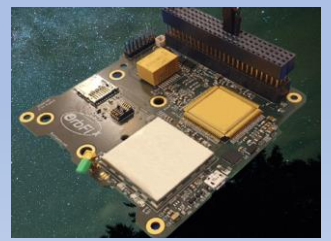

GNSS embedded system using OCEOS

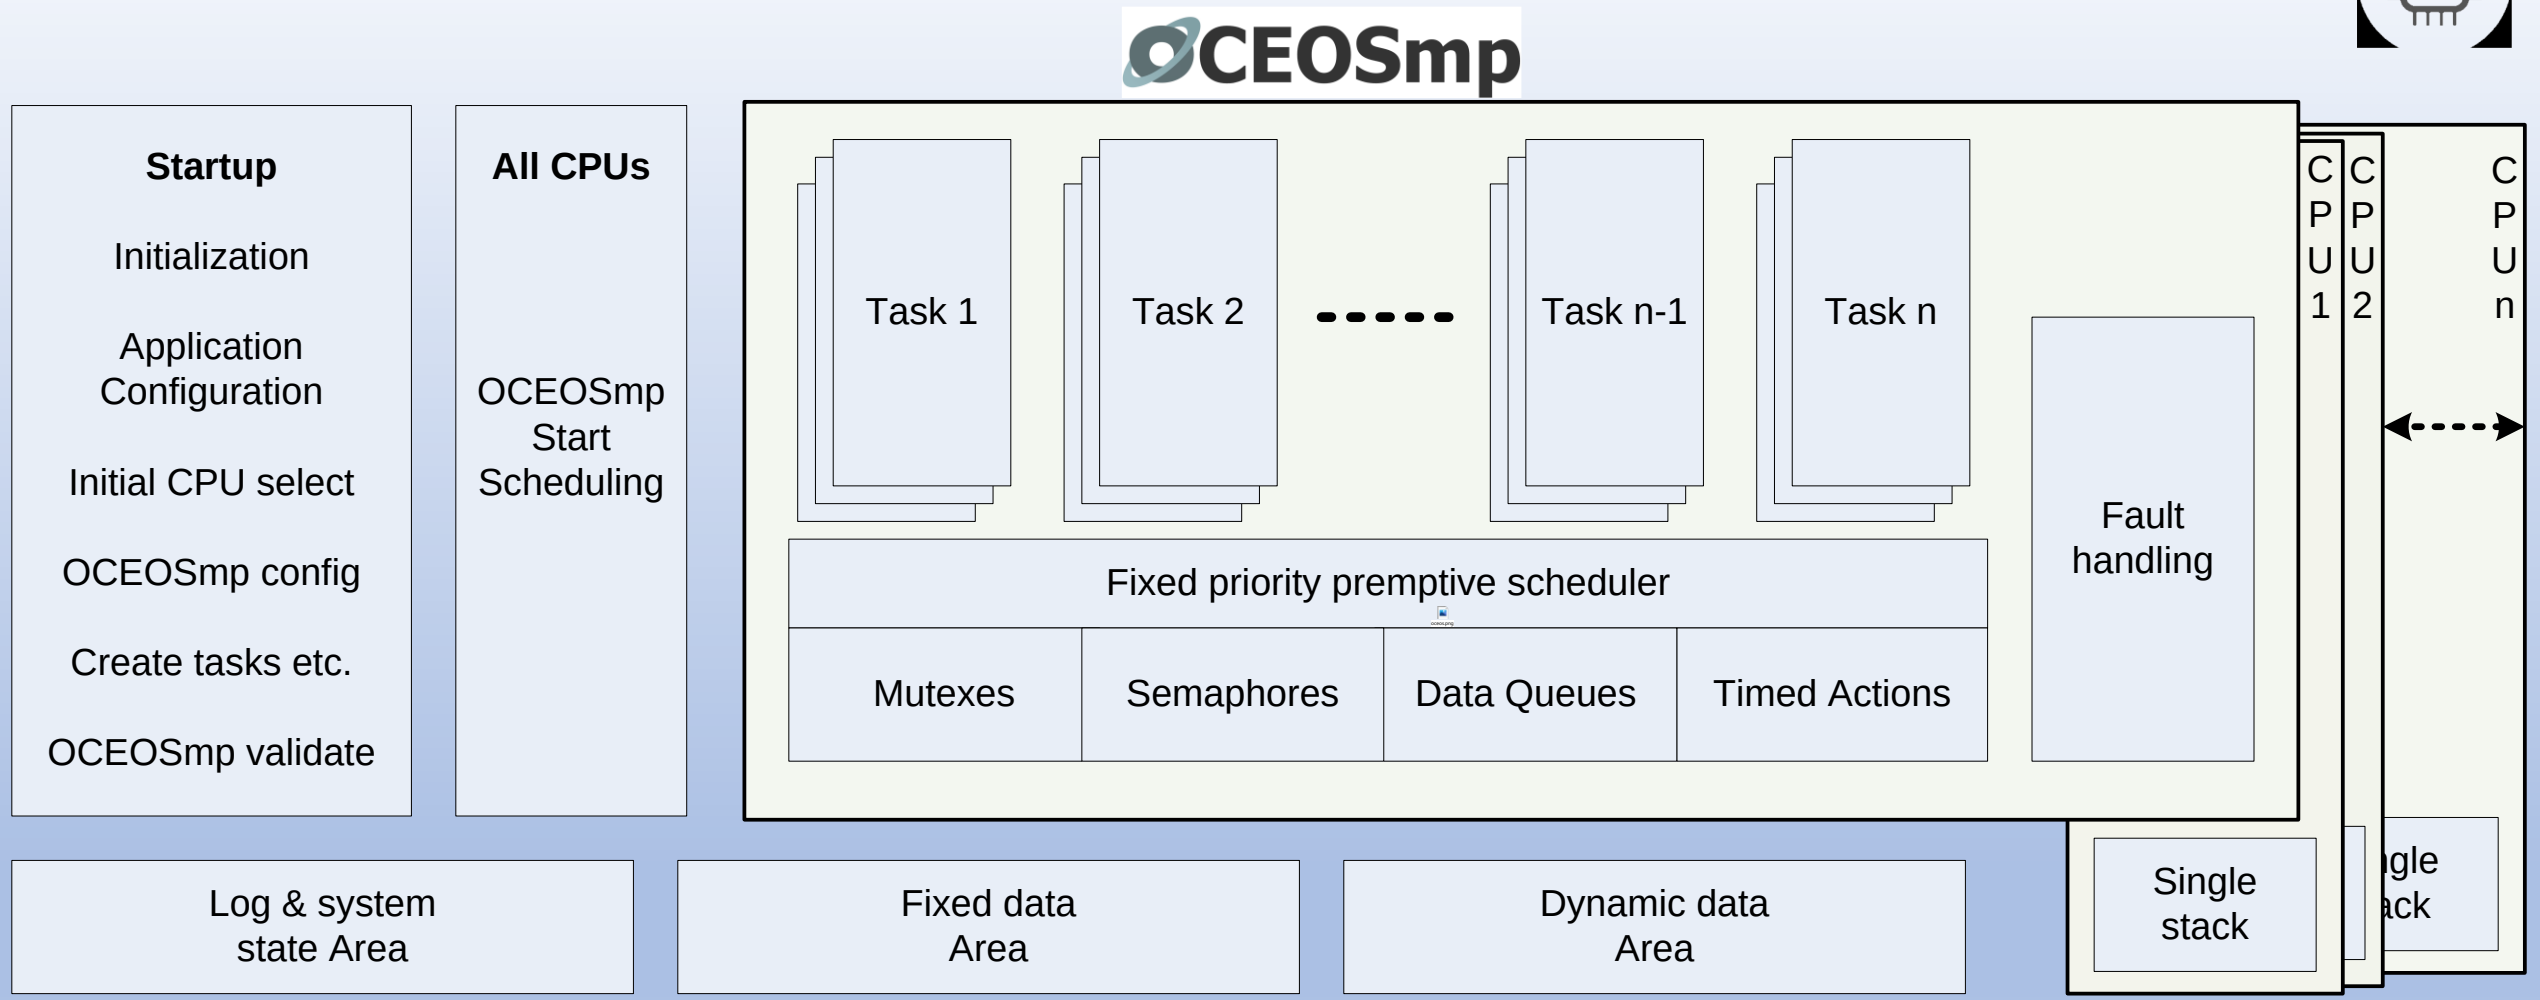
CEOSmp
Startup
Initialization
Application Configuration
Initial CPU select
OCEOSmp config
Create tasks etc.
OCEOSmp validate
All CPUs
OCEOSmp Start Scheduling
Task 1
Task 2
Task n-1
Task n
Fixed priority premptive scheduler
Mutexes
Semaphores
Data Queues
Timed Actions
Fault handling
CPU 1
CPU 2
CPU n
Single stack
Log & system state Area
Fixed data Area
Dynamic data Area

OCEOSmp: RTOS (4) – USING IT

- **Library** - components not used not linked into the executable
- **Servant not Master** – started by application main()
- **Step 1 : Create application configuration, pass to *oceos_init()***
  what cores to use, what stack space, log entries
  how many tasks, jobs per task, timed actions,
  how many mutexes, semaphores, data queues
- **Step 2: Create corresponding tasks, mutexes, etc.**
  using *oceos_task_create()* etc.
- **Step 3: Use oceos_init_finish() to complete fixed data and checksum**
- **Step 4: Pass fixed data and initial task (if any) to *oceos_start()***
  dynamic data area is set up
  multi-core scheduling begins

Debug support - DMON

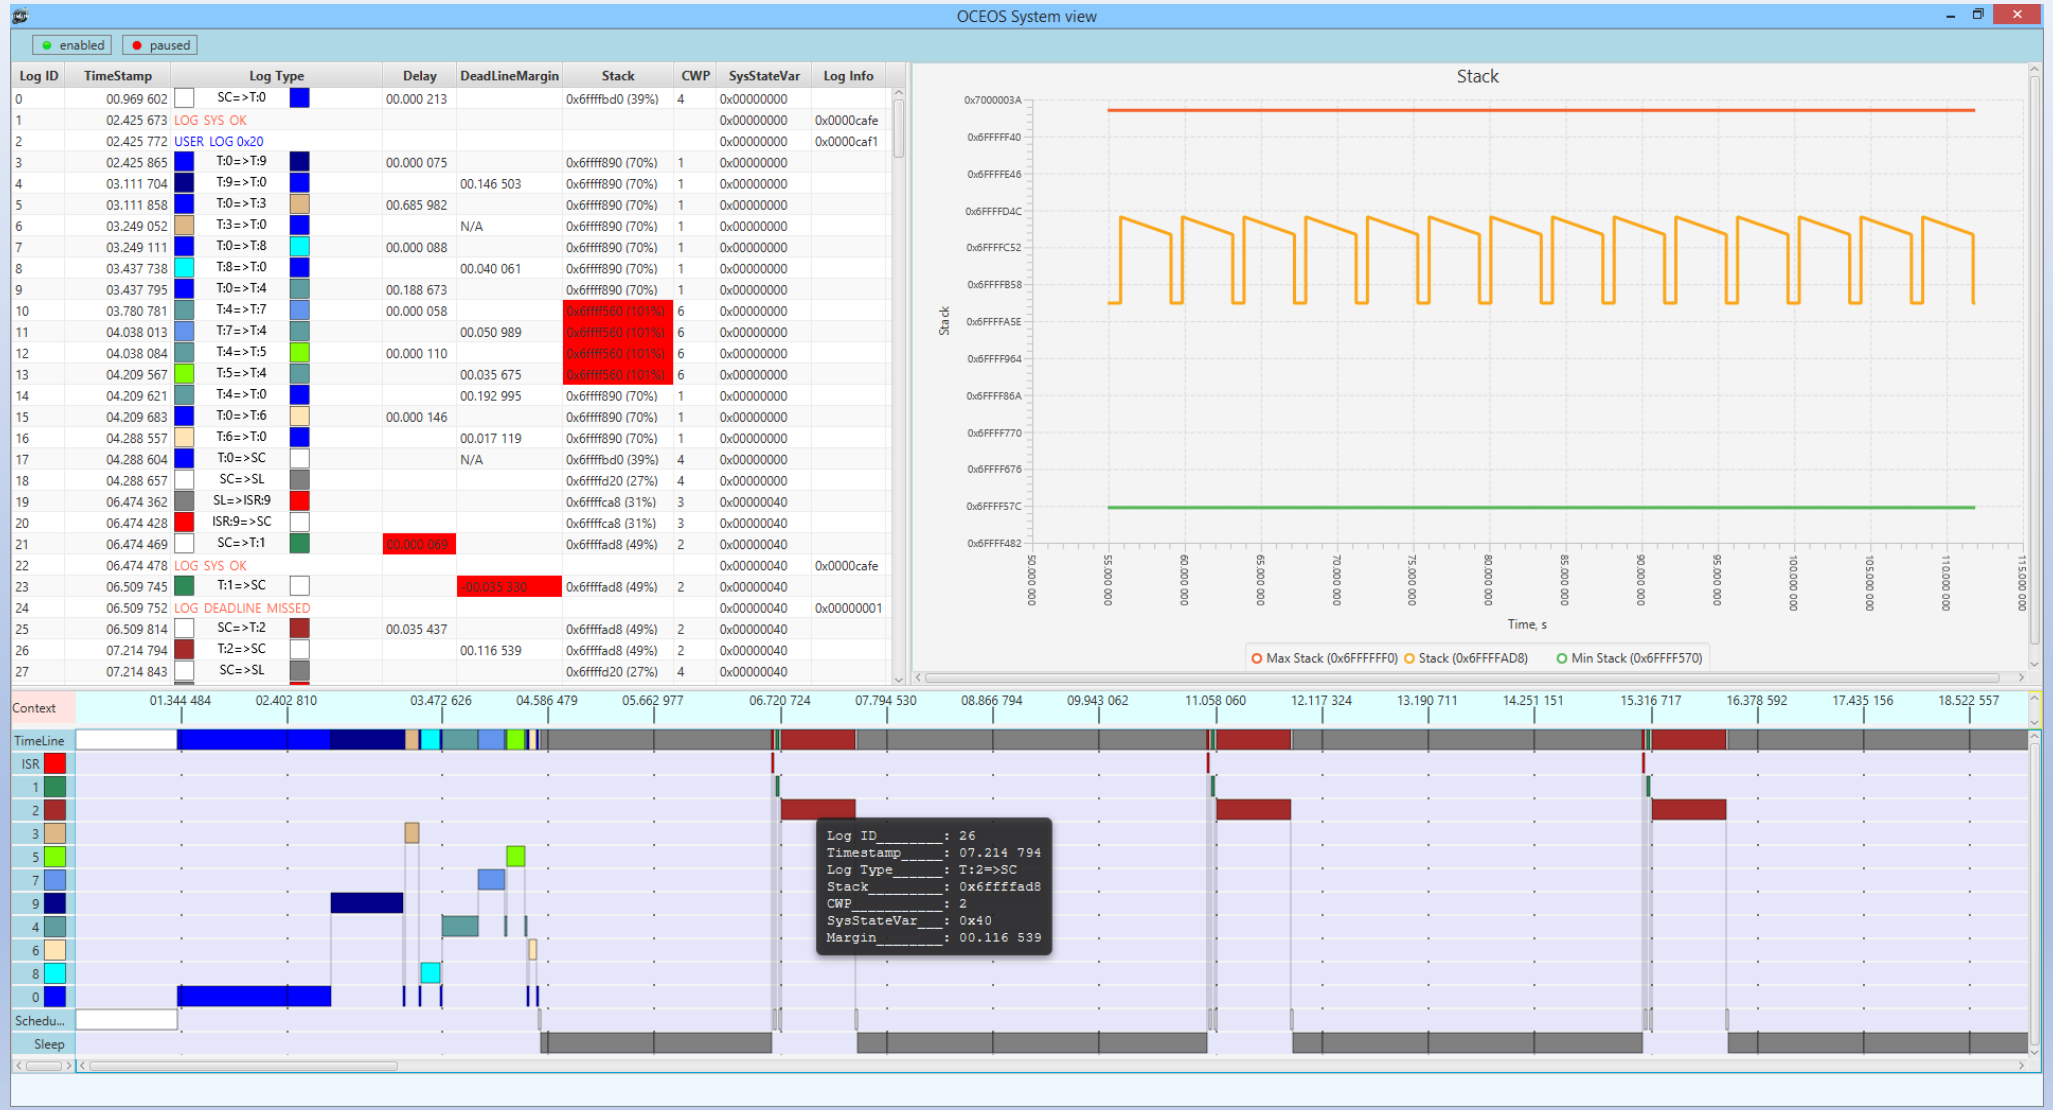

Current Status

- **OCEOS (single core)**
  - SPARC and ARM versions complete
    (with additional support for GR716 microcontroller)
    ESA Flight Level B qualification ready
- **OCEOSmp (multicore)**
  - Multicore SPARC & RISC-V scheduling in initial test, ARM later
  - Example test results using SPARC quad core Gaisler GR740:
    - 1001 task starts even distribution: Per CPU 251,250,250,250
    - 4096 sample FFT (one task, four jobs in parallel): Speedup factor 3.7
- **Availability**
  - OCEOS - single-core development kit on-sale
  - OCEOSmp - multicore beta evaluations available soon

OCEOS task usage & debug screen

Finally

- **Thanks to ESA for their support**
- **Thank you for listening**
- **Any Questions?**

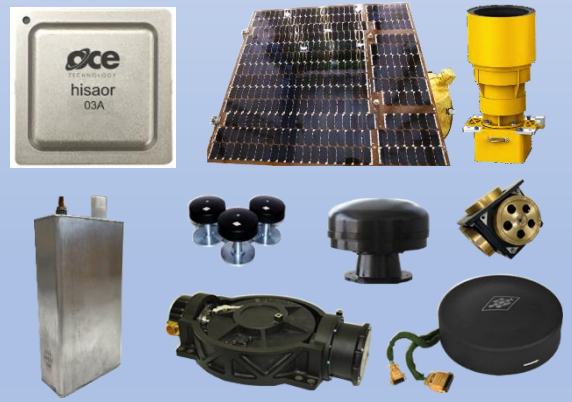

michael.ryan@ocetechnology.com

www.ocetechnology.com