

Nemesis OS Architecture

January 25, 2001

Focus of this paper: OS structure for multimedia applications (in addition to traditional apps)

MM apps:

- o temporal property: time-critical demands
- o informational property: can often tolerate some data loss

Key observation: QoS crosstalk

- o Big OS servers/services don't obey temporal properties... They mix resources among many threads/processes and have high shared overhead that affects everyone.
- o Interrupt-driven scheduling subverts the "official" scheduling policy; we can't allow interrupts to affect the overall scheduling of resources
- o All resources matter, not just CPU

Key ideas:

- o provide for dynamic alloc of resources to applications (not OS services)
- o must get the accounting of resource usage correct (eg. CPU cycles spent on paging)
- o feedback to application to adjust its resource usage (consistent vs. predictable)

CPU use managed by scheduler activations

- o threads can be resumed or can have an upcall (to the activation)
- o thread contexts managed at user level, kernel only provides atomic context switch and uses the user-level data structures
- o no kernel threads: only user-level threads with event dispatch (compare with wakeup predicates in exokernel)
- o kernel is completely non-blocking and asynchronous (if you don't have threads you can't block). This ensures that scheduling is only done at user-level (within an sdom).
- o A form of two-level scheduling

Interrupts are essentially queued as events:

- o After arrival, interrupts are masked, which limits the impact interrupts can have on system performance
- o In practice, I/O devices are polled. An (unmasked) interrupt just ensures that the driver runs to do the polling...

- o Demultiplexing is the most complicated operation a interrupt handler should do (i.e. which event queue?)

Separate control from data flow (path-based)