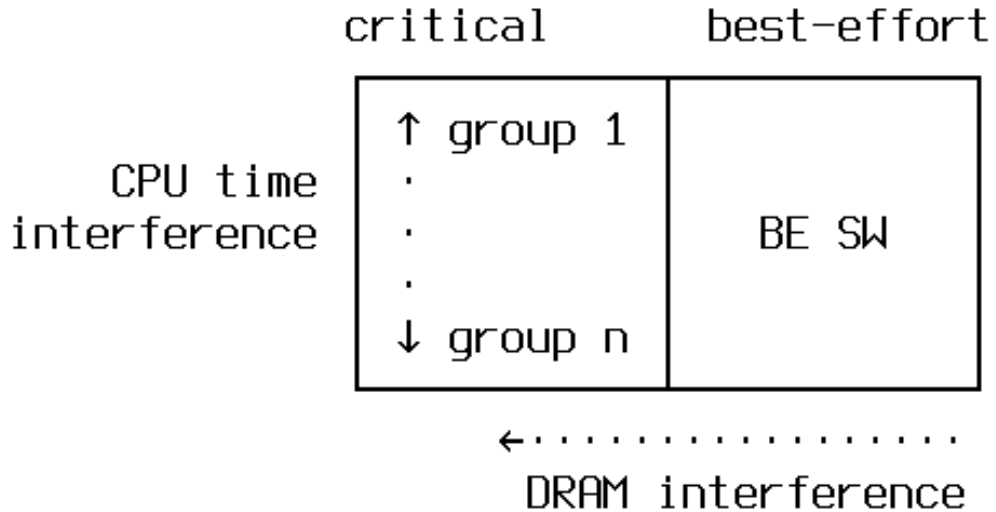


## 5.4 isolation

All critical groups must be isolated from all other critical groups, but they must also be isolated from the best-effort software. Here, the intersection is not the CPU, but the shared DRAM and the shared memory bus.



**Isolate us! Resource sharing interference.**

Isolation must be implemented in practice, as well as digested into something that is formally computable, so it can be used to determine task response times, and, later on, overall system schedulability. In this, software as well as hardware are considered.

It is important to note that the word “isolation” is not used in the general sense, but in terms of real time schedulability analysis and task response-time computation. If A is isolated from B in terms of *i*, that does not mean A is physically separated from B so that B cannot ever influence A. On the contrary, it means that A *is* influenced by B precisely through *i*, only the designers and implementors of the system has accounted for this, for example by providing throttling mechanisms that in practice will provide a bound for the worst-case interference B can have on A (through *i*). This bound is then used in the response-time calculations, which is the theory mirror image of the real-time implementation.