Unified Memory Behavior

When **UM** is allocated, it may not be resident initially on the CPU or the GPU

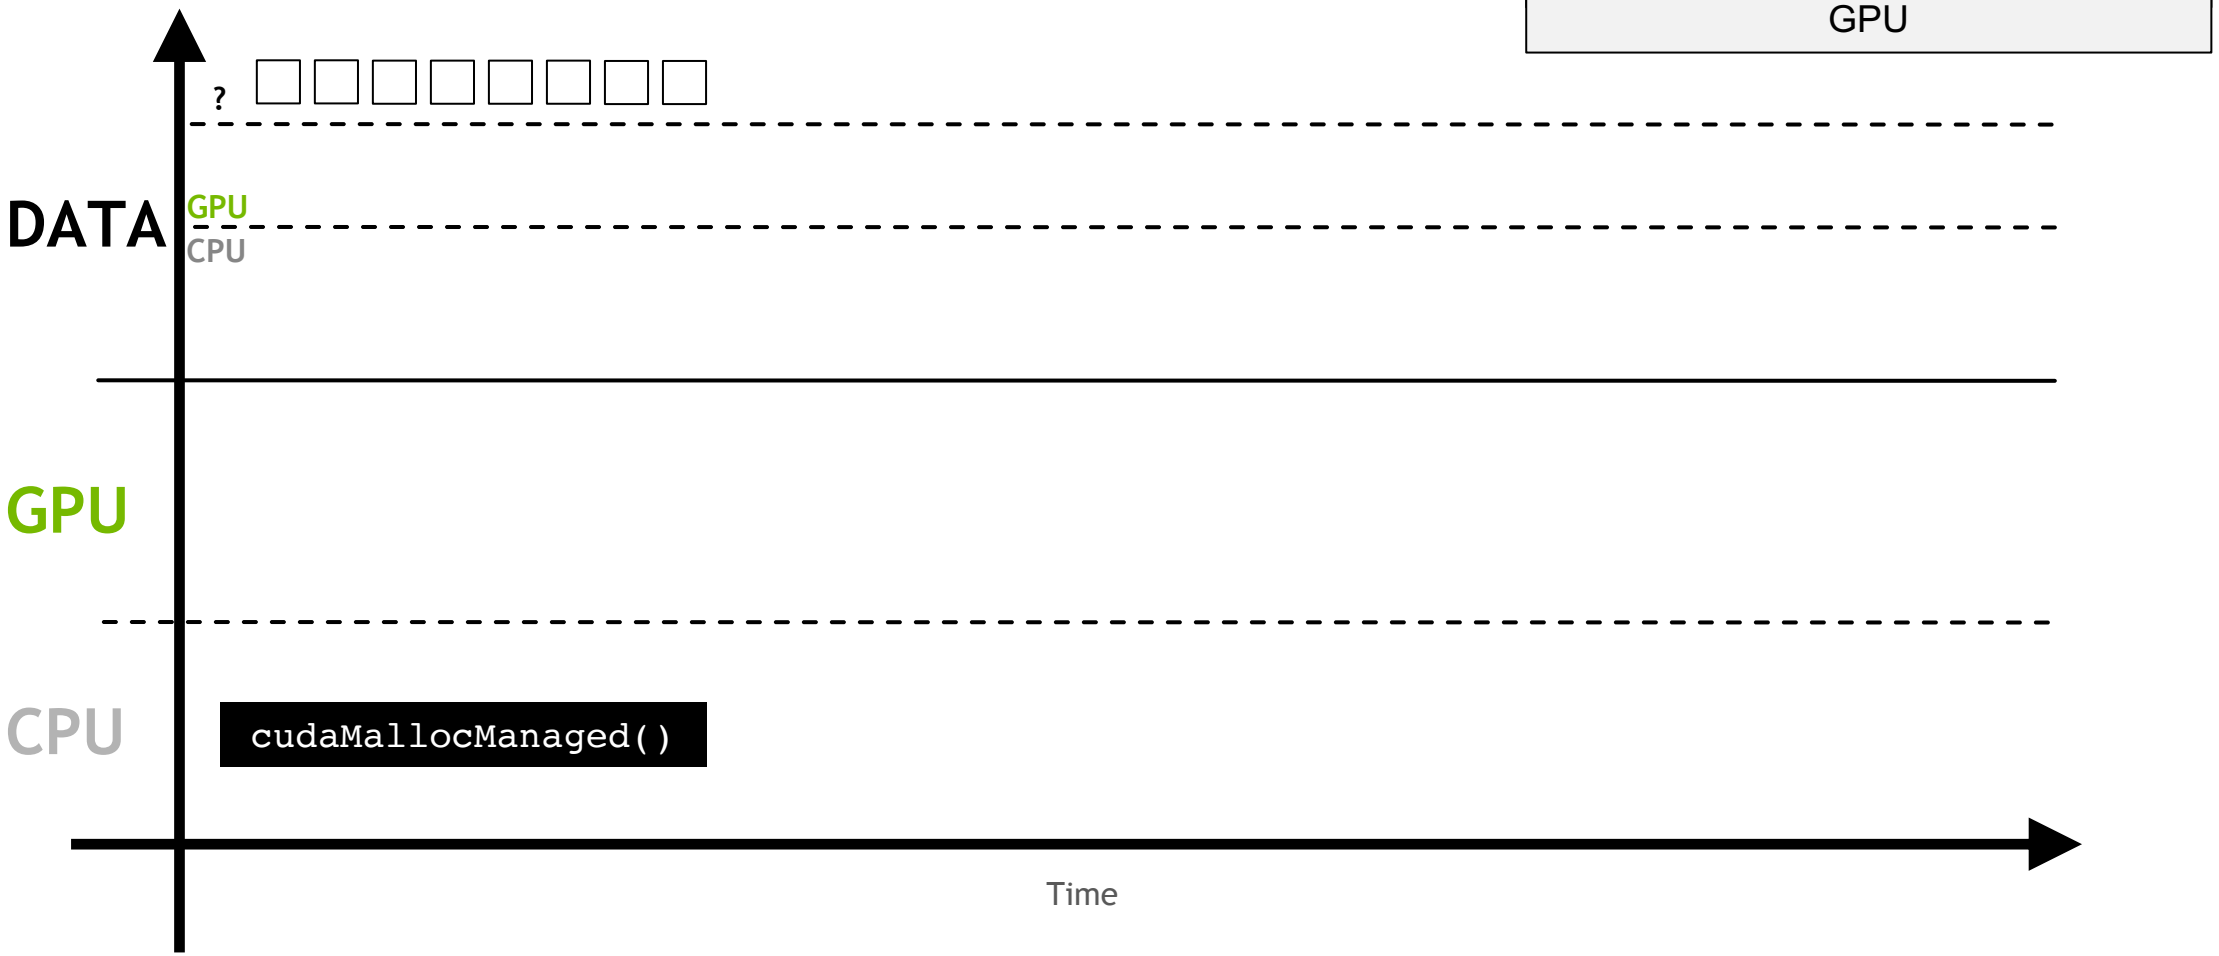

When some work asks for the memory for the first time, a **page fault** will occur

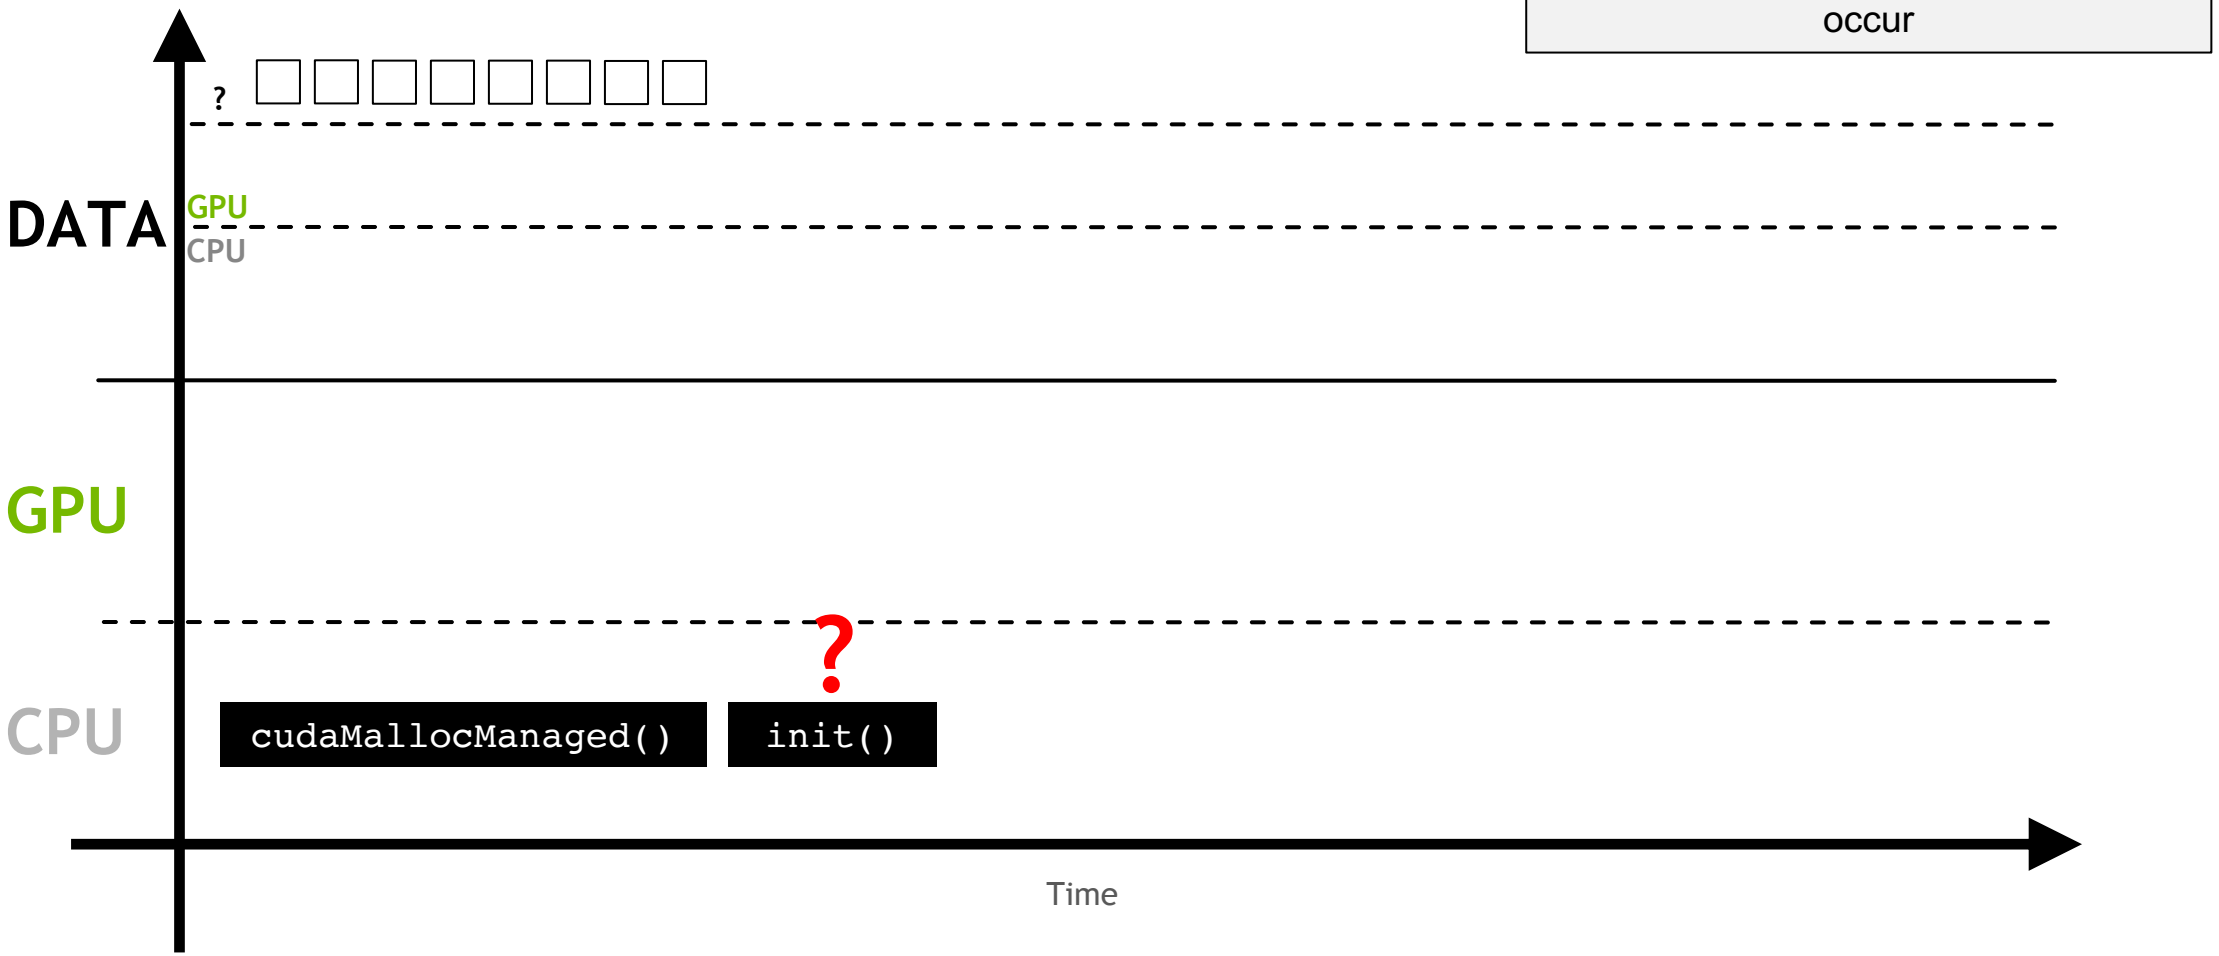

The page fault will trigger the migration of the demanded memory

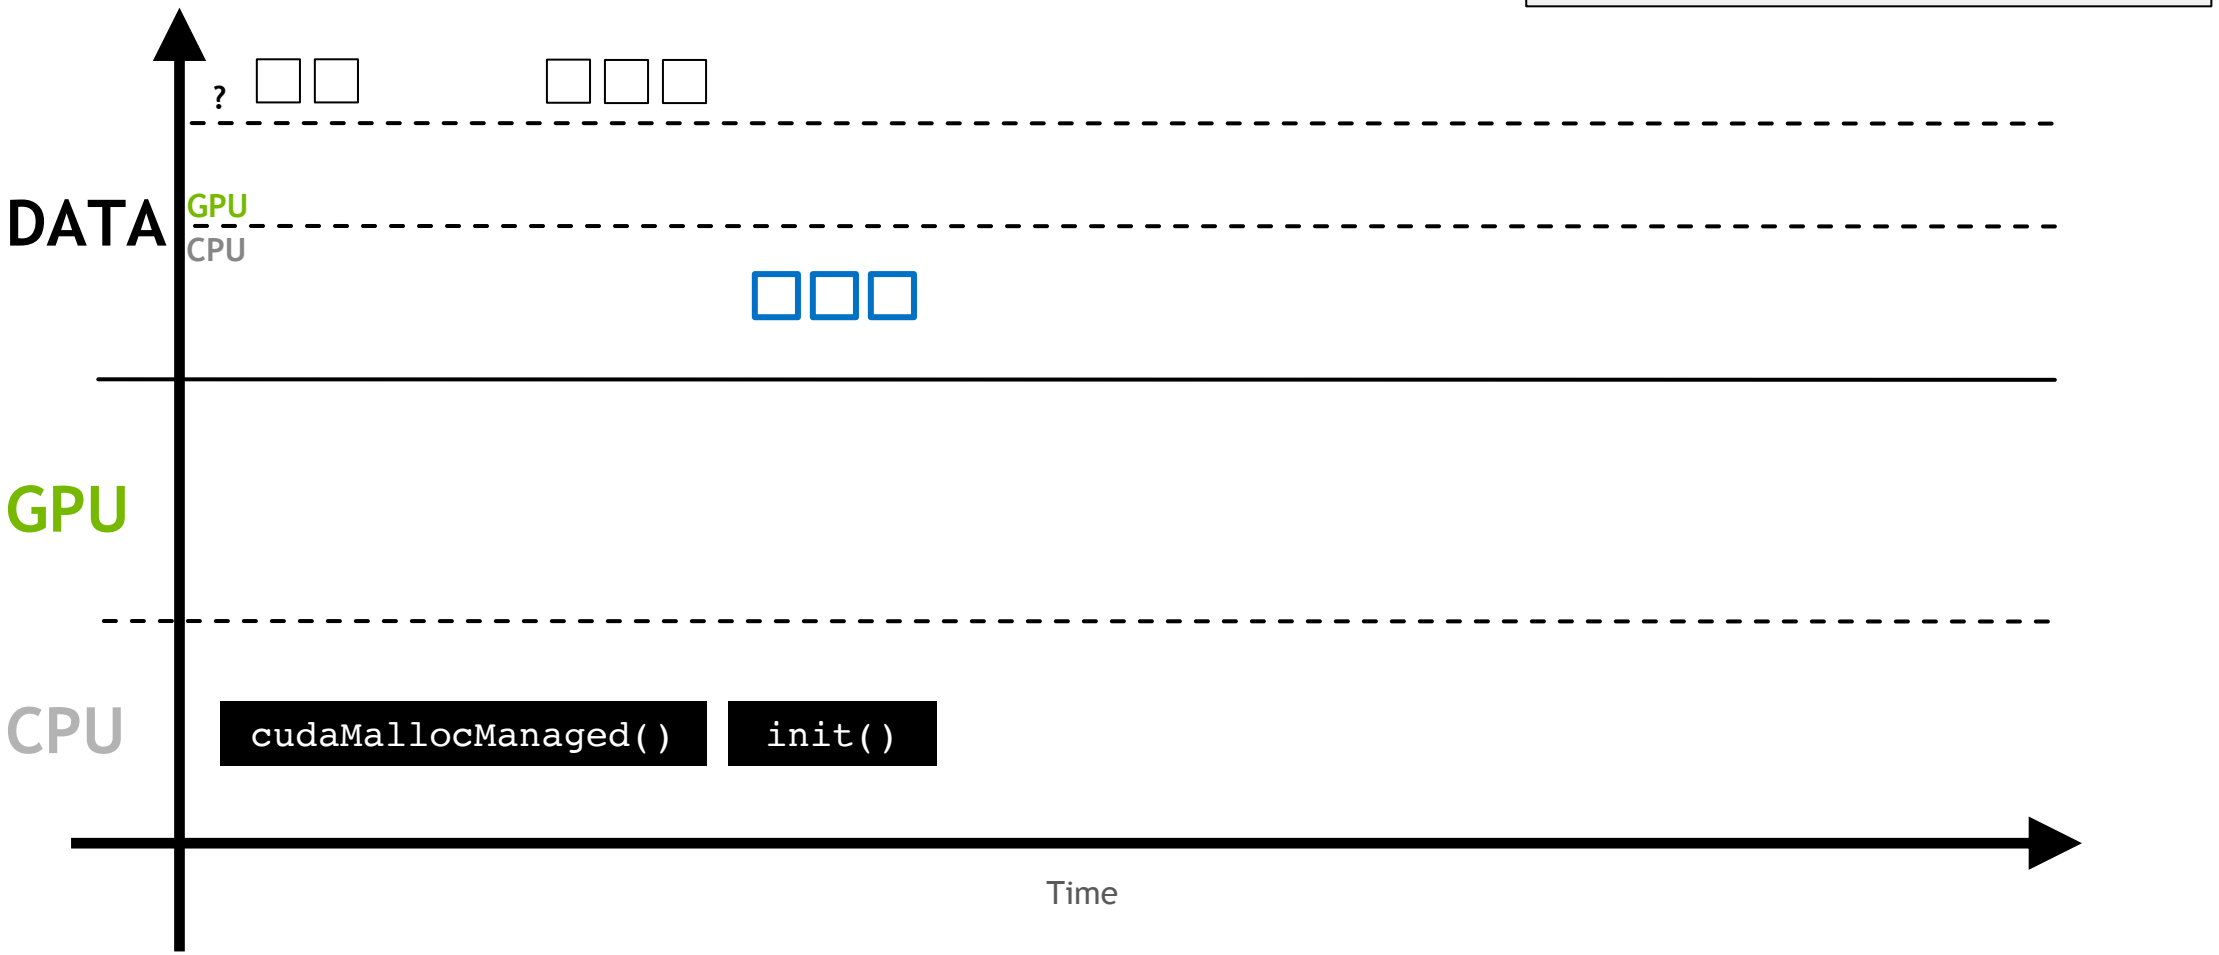

This process repeats anytime the memory is requested somewhere in the system where it is not resident

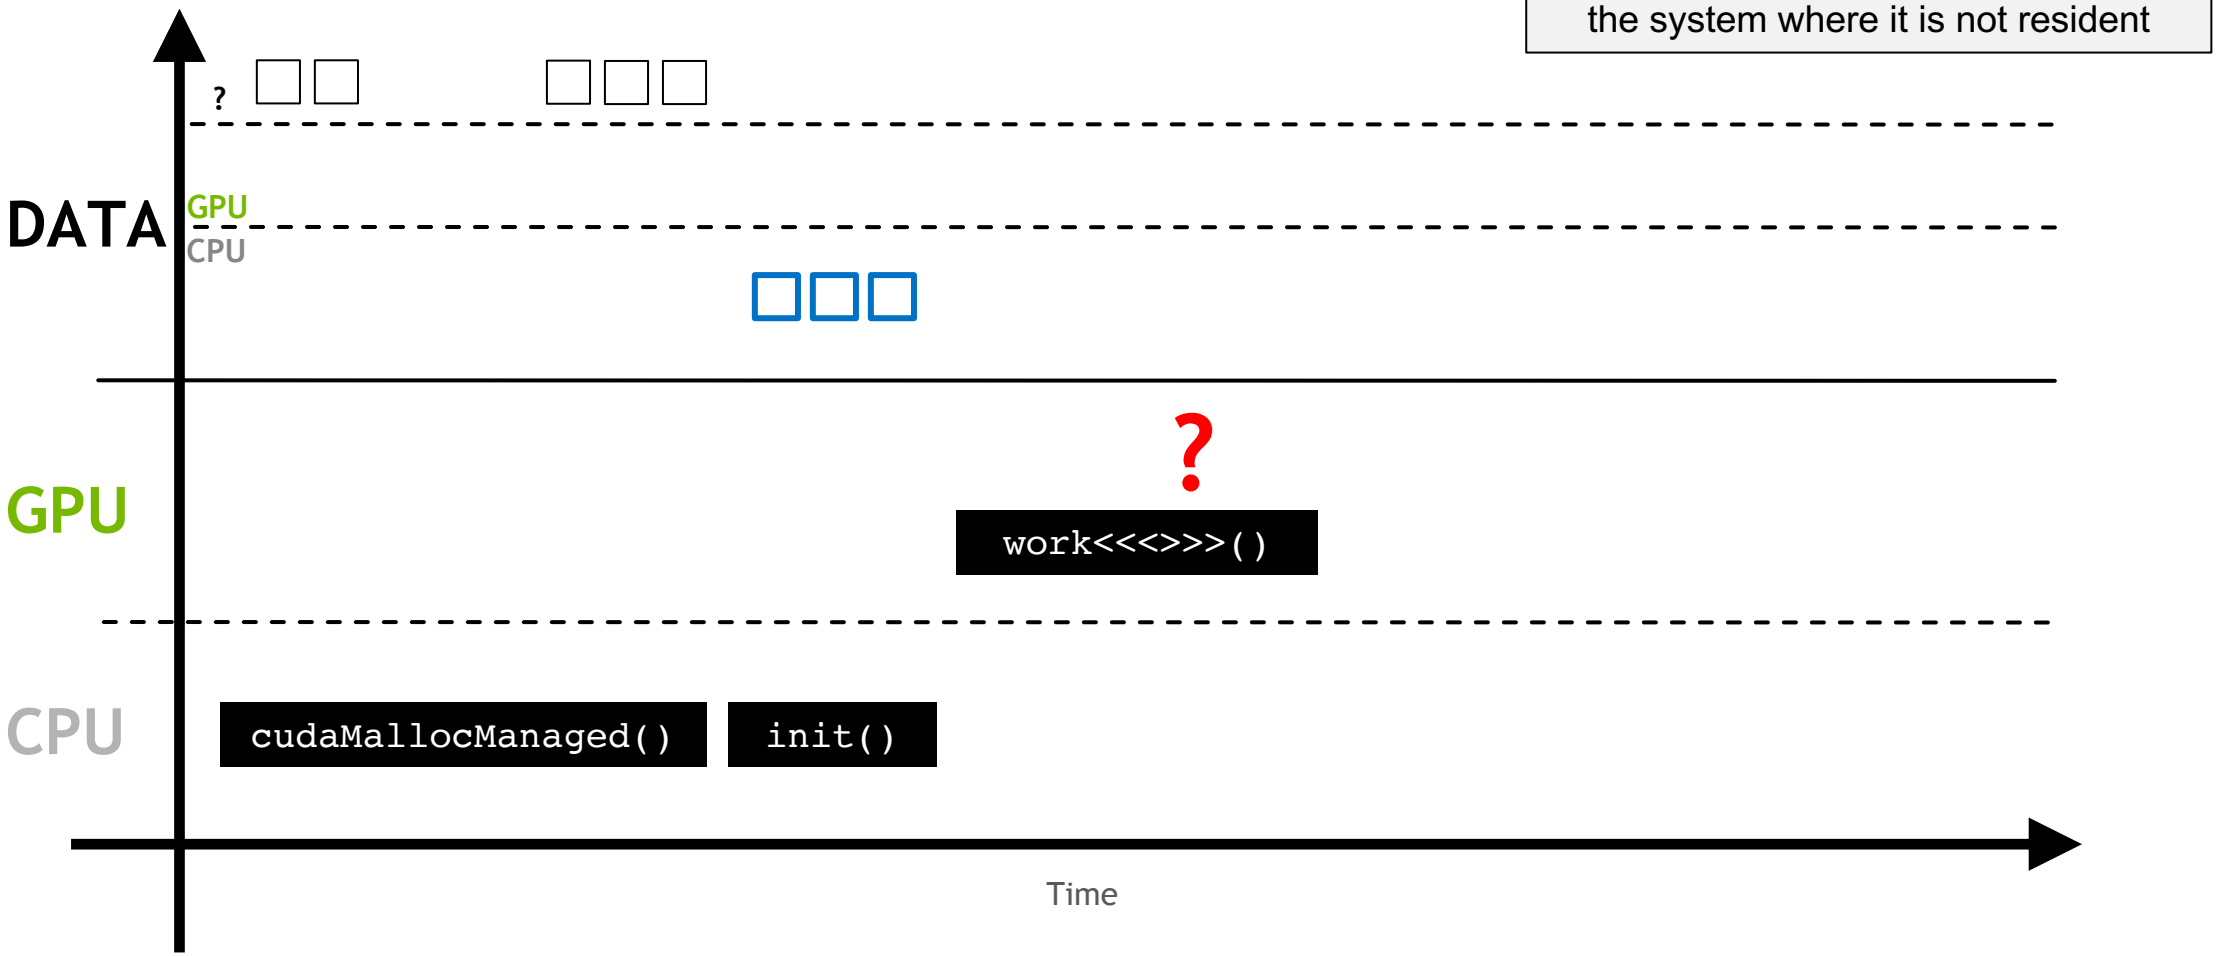

This process repeats anytime the memory is requested somewhere in the system where it is not resident

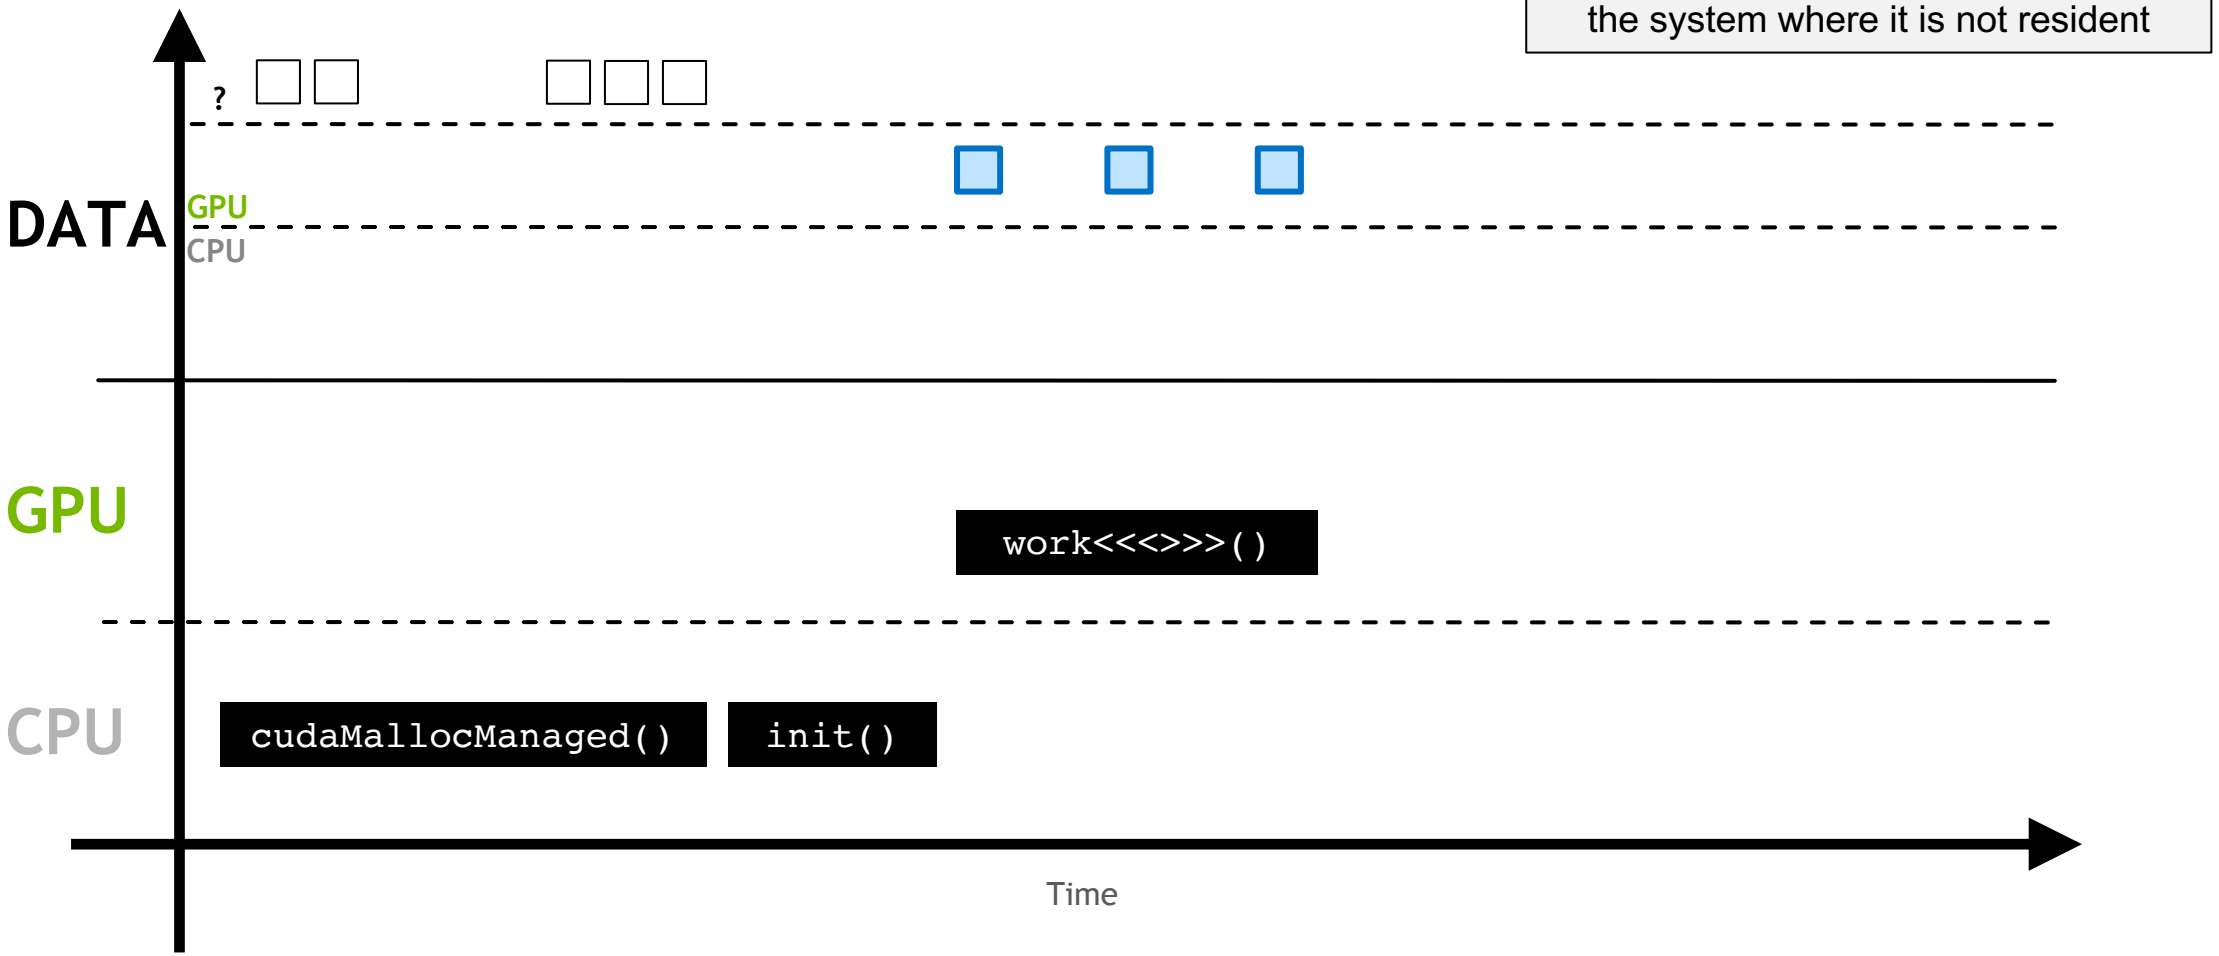

If it is known that the memory **will be** accessed somewhere it is not resident, asynchronous prefetching can be used

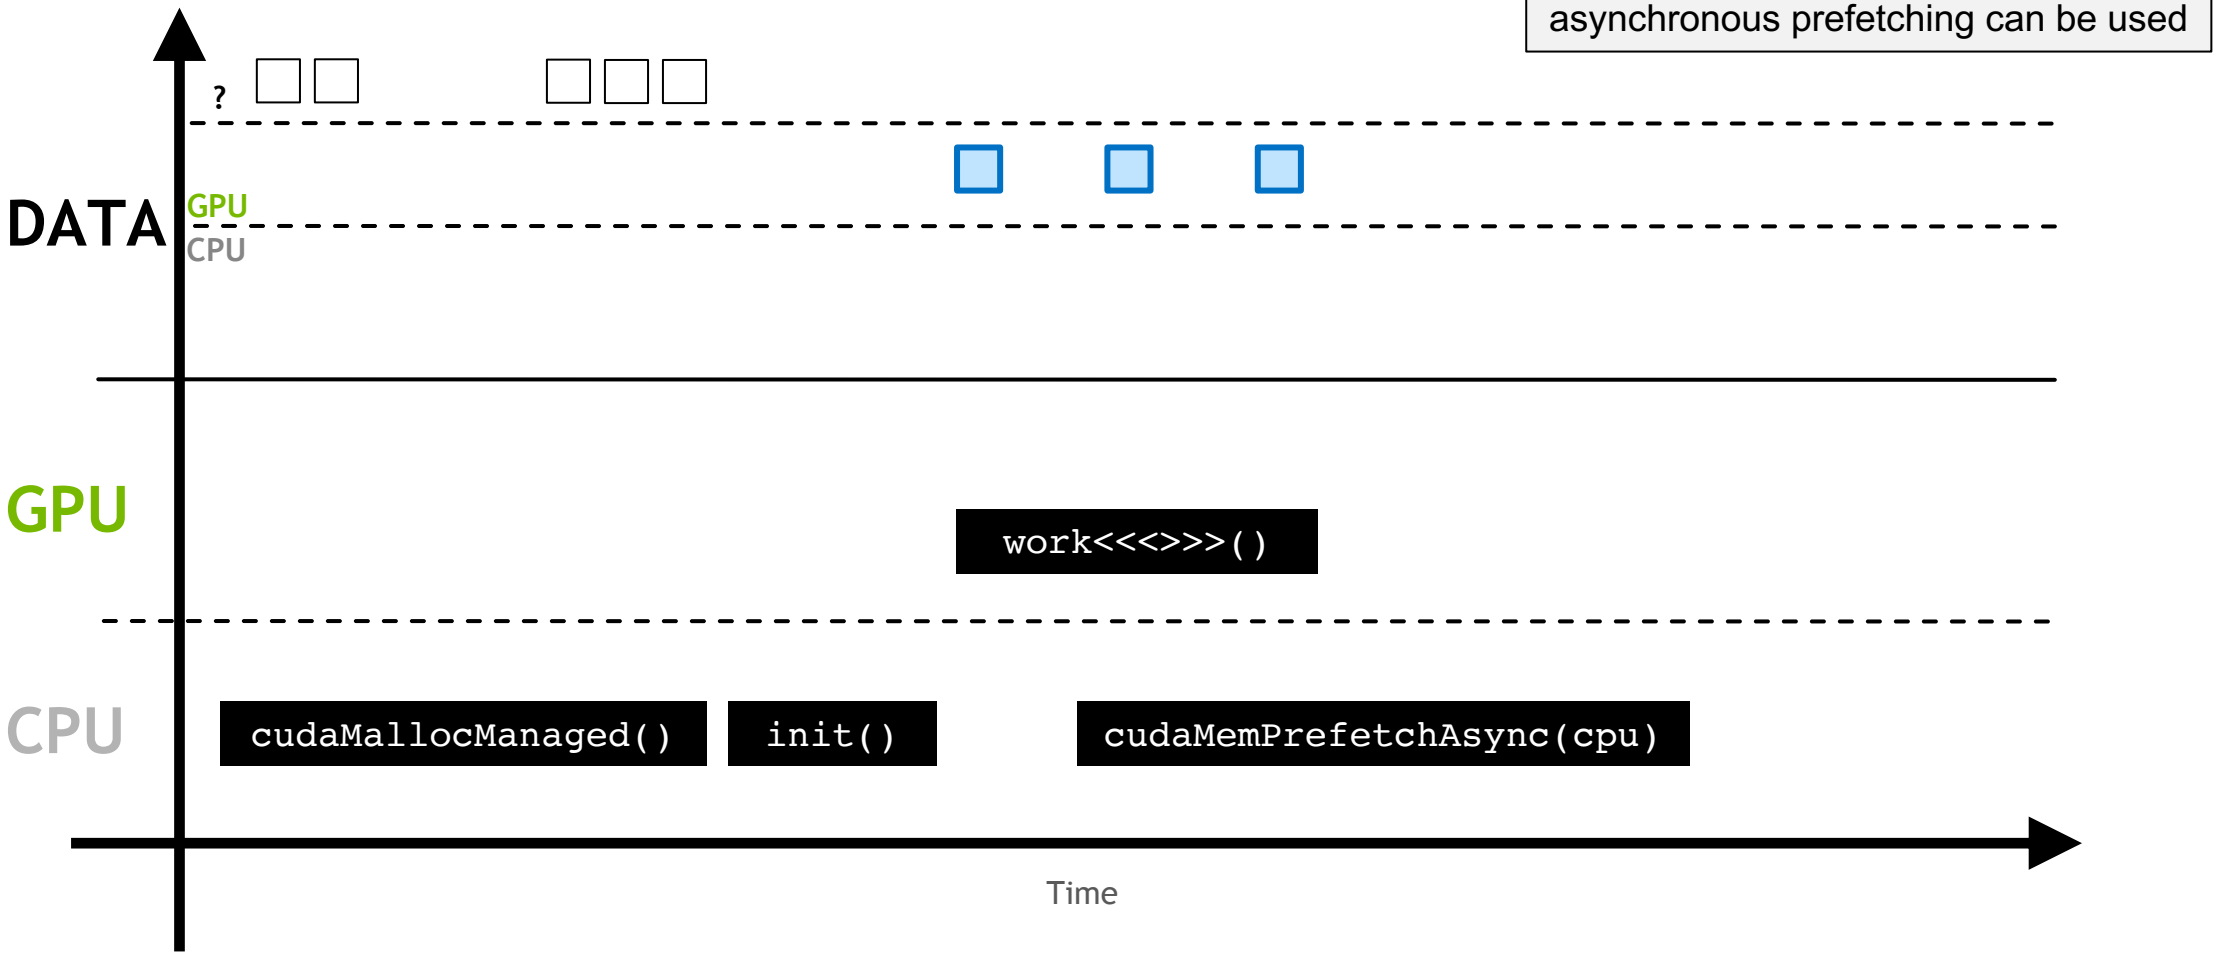

This moves the memory in larger batches, and prevents page faulting

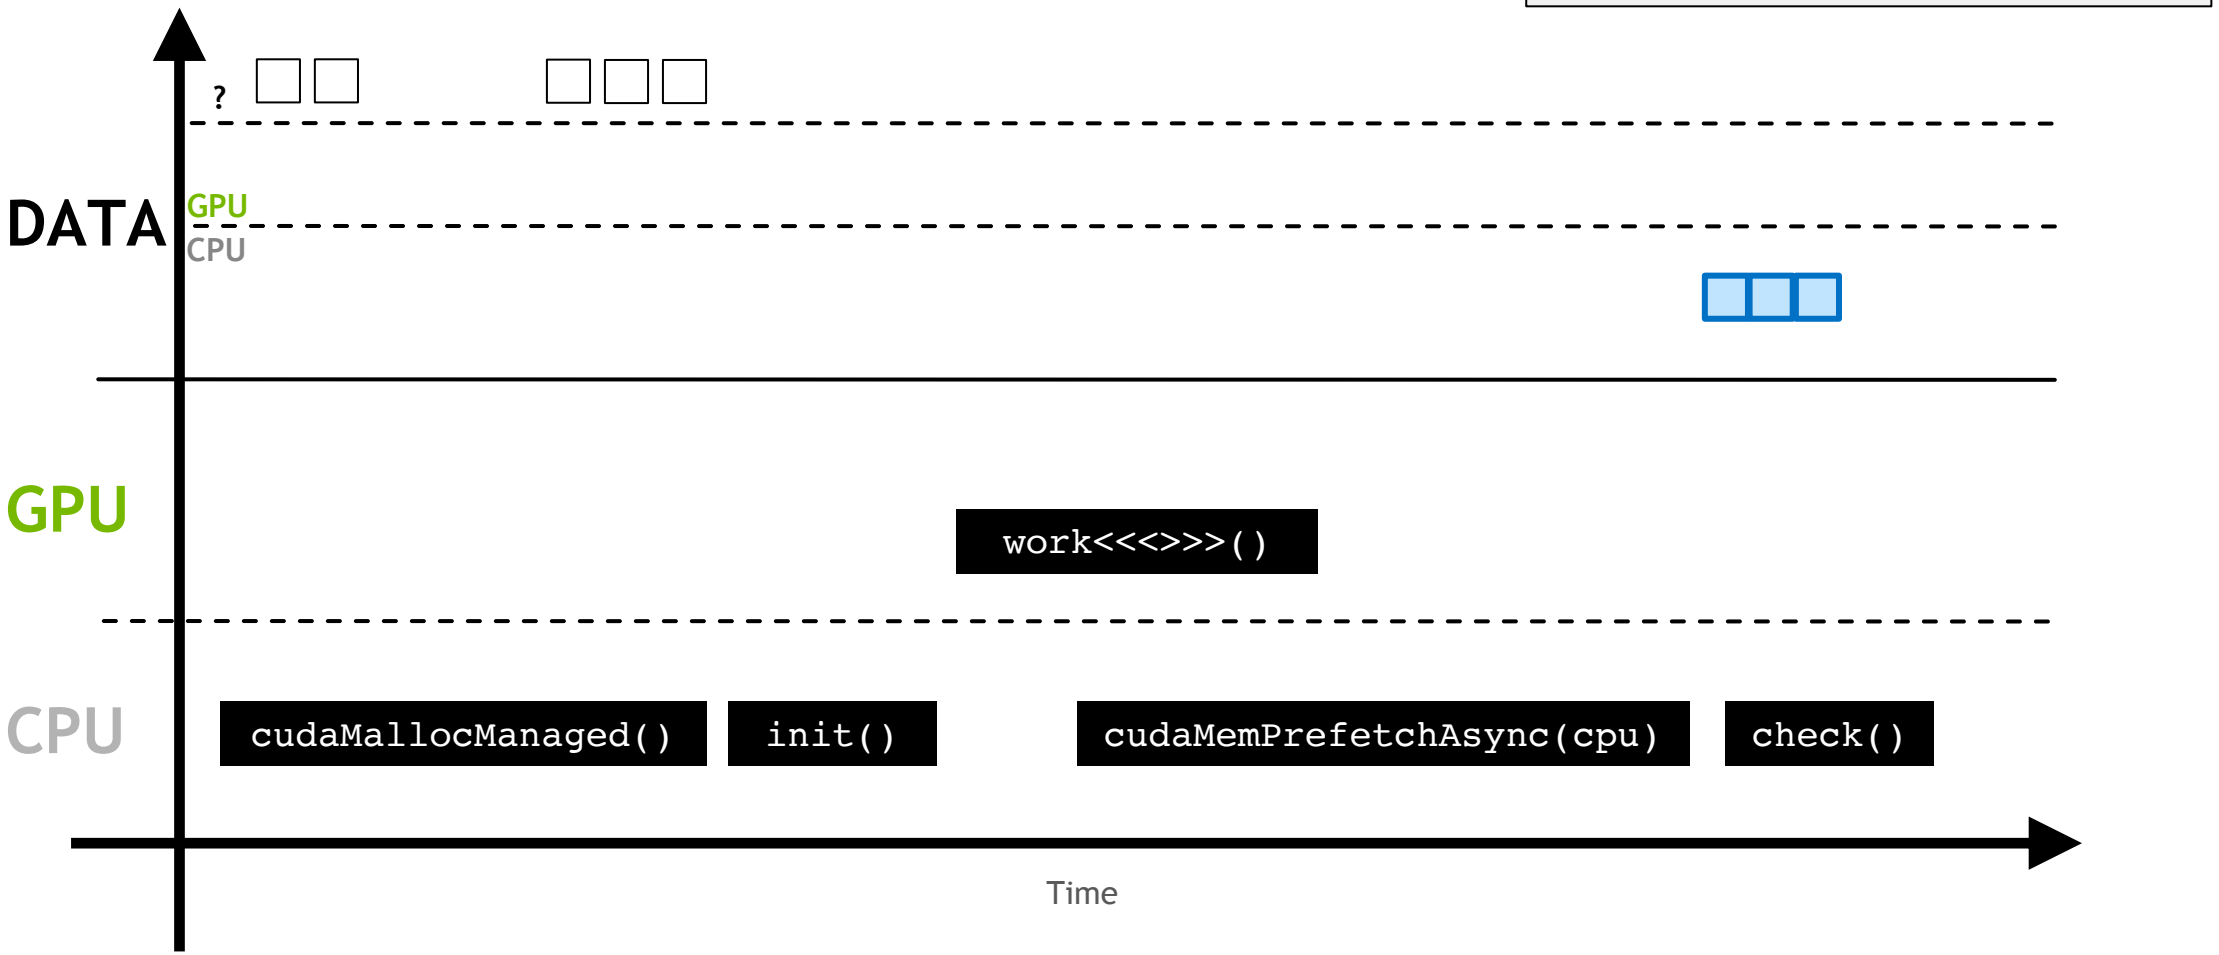

www.nvidia.com/dli