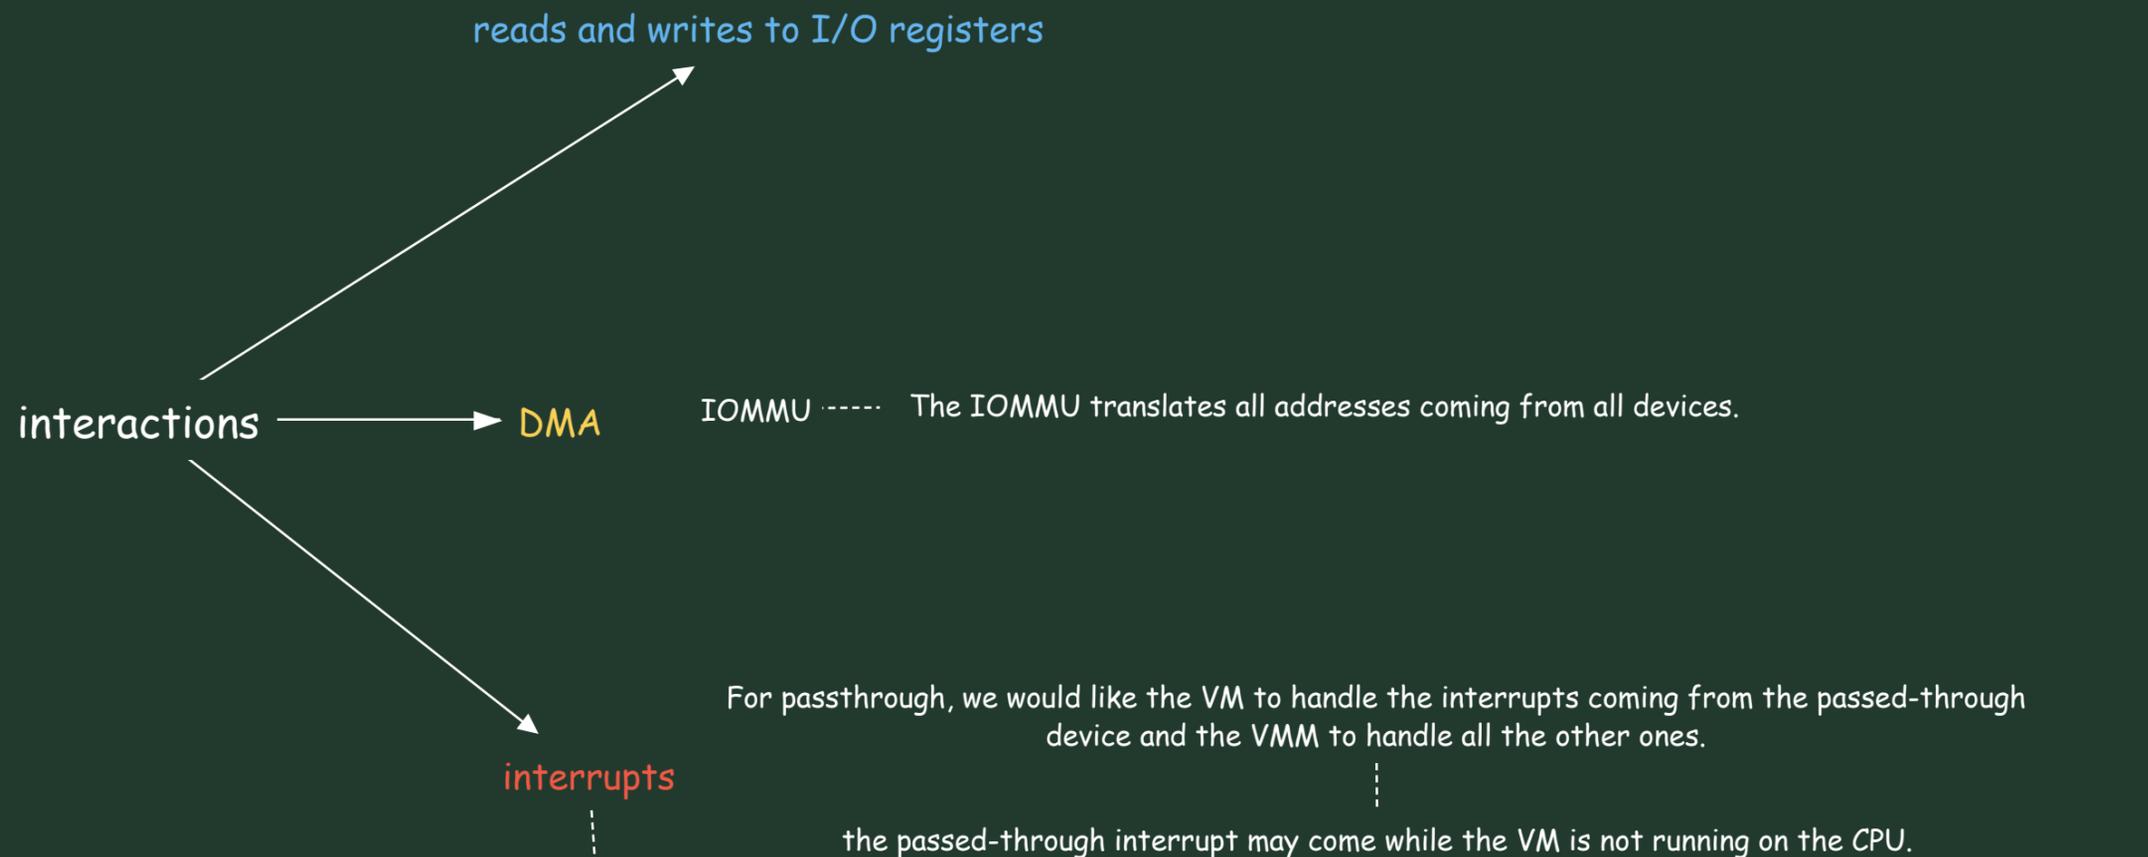


PCI SR-IOV

----- This is a feature of some PCI devices that may offer several virtual instances of themselves. The instances may be used independently, even if they internally share some hardware.

To implement passthrough in hardware assisted virtualization, we need help from hardware.

----- This is because the VMM software is not running, in general, when the interactions between the guest and the peripheral take place, and therefore the VMM needs help from the hardware to guarantee that these interactions are properly handled.



For passthrough, we would like the VM to handle the interrupts coming from the passed-through device and the VMM to handle all the other ones.

the passed-through interrupt may come while the VM is not running on the CPU.

Now let's assume that a device D that has been passed-through to VM M generates an (host) interrupt. This is what we would like to happen:

1. if the M CPU is Running, we want the interrupt to be handled by the processor, jumping to the guest interrupt handler, without any VM exit (and, therefore, without any intervention of the VMM);
2. if the M CPU is Ready, the host CPU is currently running something else, e.g., another VM CPU; in this case we want to store the interrupt request somewhere, so that the M CPU may handle it at the later time when the M VM is scheduled and becomes Running;
3. if the M CPU is Halted, we need to store the interrupt request, as above, and also make the M CPU become Ready.

← posted interrupt mechanism