Virtual Reality - Quiz #3

Solutions

1. **Draw and explain the functioning of the CyberForce. What does it have in common with the Phantom and the CyberGrasp interfaces? Can it produce torque feedback at the wrist? Why yes or no? (2 points)**

   The CyberForce is a combination of a CyberGrasp force feedback glove with a serial-kinematic force feedback arm attached to the wrist. The arm has six embedded sensors, and serves as a tracker, providing the wrist position to the Force Control unit. Finger position information is provided by the CyberGlove component of the CeberGrasp system. Once the position of the user's wrist and fingers is sent to the host computer running the VR simulations, its processors detects collisions and calculates resulting forces. These data are sent to the Force Control Unit, which in turn sends commands to the CyberForce and CyberGrasp Actuator units. This way the user has force feedback at the tip of the fingers during flexing, as well as translational forces at the wrist. Since the CyberForce Arm has only three actuators, it cannot produce torque feedback at the wrist. Furthermore, the dimensions of the arm limits the user's freedom of motion, as the arm is attached to the desk.

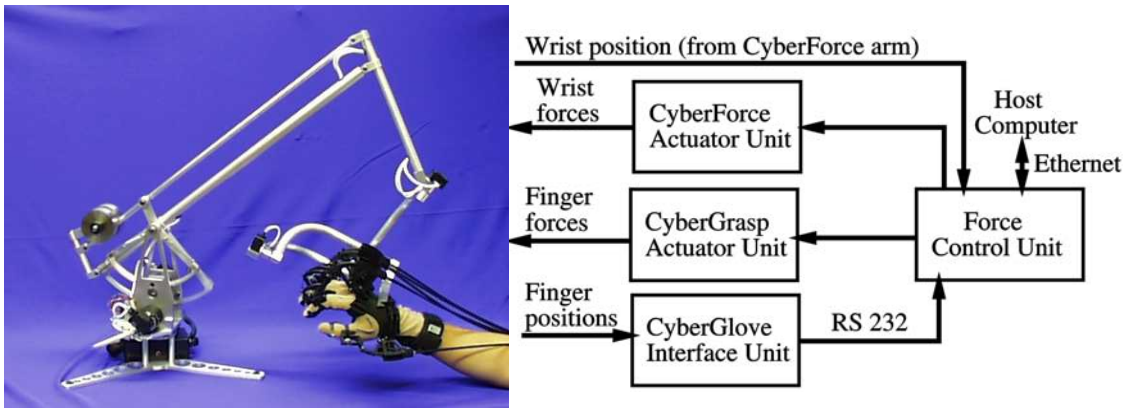

2. **Why are multi-core architectures important in VR? How is the multi-core architecture implemented in the PlayStation 3? Draw and explain (2 points)**

   Multi-cores architectures are important in VR due to the requirement that computations be done in real time. Such computations relate to collision detection, haptic feedback (physical modeling), object behavior, 3-D sound and user I/O. Typically the large computational load requires multi-processor PCs, and computation parallelism using hardware accelerators. Multi-core architectures will obviate the need for separate hardware resources such as 3-D cards, or PPUs, since computation can be done by the cores. Communication between cores will always be faster than communication with outside resources over the motherboard bus.

   The PlayStation3 processor (made by IBM) has one PowerPC managing processor and eight cores (SPE 0…7) interconnected on an Element Interconnect Bus.

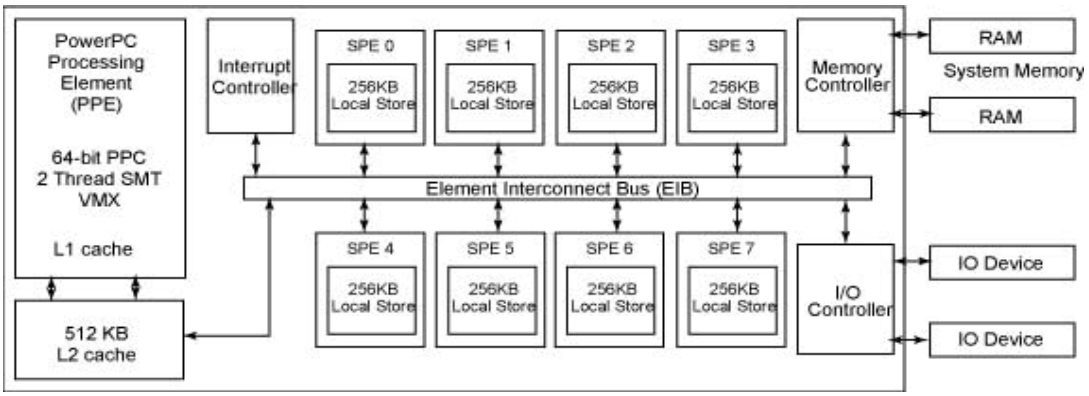

3. **If you had two side-by-side monoscopic displays, and one Phantom haptic interface, how would you synchronize the graphics and the haptics pipelines? Assume the displays are driven by separate PCs. (2 points)**

   The two PCs synchronize their graphics pipelines using a multi-view cable, which makes one of the PC the master, and the other the slave. When the slave PC graphics pipeline is finished filling the output buffer and is ready to swap, it lets the master PC know, and the master PC raises the release signal when it finished its graphics computations. Both PCs then swap the buffers to their respective displays. Since only one Phantom haptic interface exist, it needs to be connected to the Master PC. In this way the haptic pipeline application stage is synchronized with the Master PC graphics pipeline application phase.

4. **Now assume that your simulation is displayed on a wall-type tiled display with back projectors. Where is the bottleneck on the numbers of tiles you can have? Give one way to address this limitation. Which is a better approach to driving the projectors – PC clusters or rack of blades and why? (3 point)**
   We assume that each projector (and tile) has its own graphics pipeline feeding it. These graphics pipelines can reside in separate PCs (in case of PC clusters), or be blades in a rack-mounted cluster. The bottleneck on the number of projectors (thus tiles) is the LAN interconnecting the Managing PC with the PCs driving the projectors. This LAN has a limited bits/sec throughput capacity, and will saturate when too many PCs are connected over it. To have many tiles and still maintain good frame rates (compared to having a single projector), the communication needs to limit the number of bits. Two approaches is a) update only the PC driving the projector where there is a change in the virtual scene; b) use encoding such that the messages sent are smaller in size. The rack of blades is the better approach as it is more economical in terms of space (for the same number of graphics pipelines) than an equivalent PC cluster.

5. **How can you have peers of users on "both sides of the pond" (Atlantic) interact in the same simulation? Draw and explain (1 point)**
   Communication between peers at long distance can only be done over the wide area network (WAN). Beers however use broadcast communications, which are not allowed over WANs (due to the danger of saturating the net). In order to solve this contradiction, each group of peers needs a proxi server. Each of the two proxi server encapsulates multicast packets from its peers into unicast ones, and sends these

unicast packets to the other proxi server. Each proxi server also unwraps incoming packets and then broadcasts these over its LAN.