Example Shadow Register

Freely submitted by IBM for OASIS DITA Semiconductor Information Design Subcommittee.
There is nothing proprietary or confidential about this example, and it can be shared with SPIRIT or IP-XACT as needed for discussion.

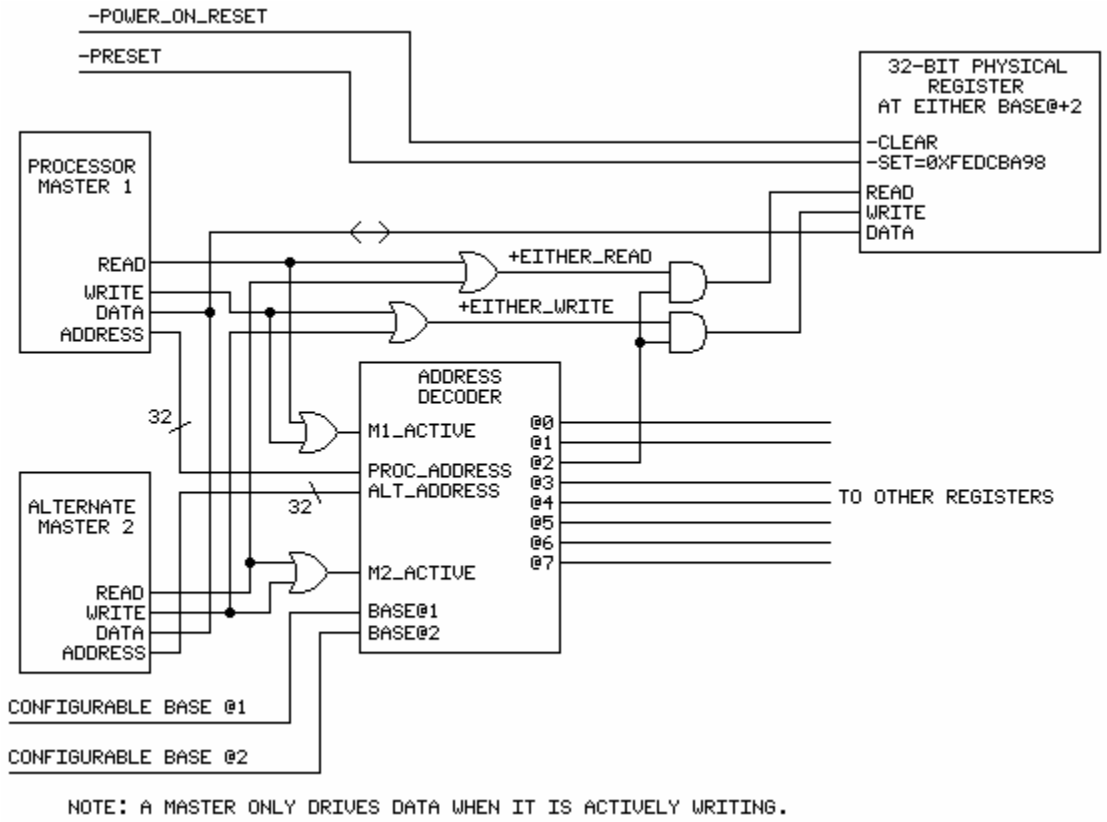

Operation: Two masters can read or write this single physical register. The processor master 1 can access any of the eight decodes from CONFIGURABLE BASE @1 + 0 through CONFIGURABLE BASE @1 + 7. The alternate master 2 can access any of the eight decodes from CONFIGURABLE BASE @2 + 0 through CONFIGURABLE BASE @2 + 7. Either of the base addresses can be even or odd. For simplicity's sake, let's not consider the accesses to be byte wide, but only one full 32-bit word at a time, so that the least-significant address bit actually addresses a full 32-bit word address.

M1_ACTIVE and M2_ACTIVE determine which master gets priority access to the address decoder. M1_ACTIVE gets priority. The diagram does not show access collision control between the masters.

So, as an example, consider if configurable base @1 is set to 01234567h and configurable base @2 is set to 12345678h.

The one physical register is readable and writable by both masters, but through very different address spaces.
The address for the processor master 1 would be 01234567h + 2 = 01234569h.
The address for the alternate master 2 would be 12345678h + 2 = 1234567Ah.

Note that the CONFIGURABLE BASE addresses could be programmable by the user through some other configuration register or it could be hard wired by the designers. Also note that this example physical register can be cleared to all zeros by the signal –POWER_ON_RESET, and can be preset to a value FEDCBA98h by the assertion of -PRESET.

I think the important thing to understand is that there is really no limitation to the number of different ways the register could be cleared, preset, or addressed through different masters.