CLAIMS:

1. A heterogeneous three-dimensional circuit stack comprising:

a first processor;

a second processor communicatively coupled with the first processor through one or more through silicon vias, wherein the first processor and the second processor are graphics processors; and

interconnect logic to communicatively couple the first processor and the second processor to a shared resource, the interconnect logic to couple with the first processor and the second processor via corresponding on-chip interconnects;

wherein the interconnect logic includes bandwidth sharing logic to adjust bandwidth to the shared resource;

wherein the shared resource includes memory to cache data to be received via the interconnect logic;

wherein the interconnect logic is to operate at a higher frequency than one of the first processor or the second processor; and

wherein the interconnect logic is to divide a message transmitted from the first processor or the second processor to another processor into a number of separate messages based on a frequency ratio of the interconnect logic, the first processor, the second processor and the another processor.

2. The heterogeneous three-dimensional circuit stack as in claim 1, wherein the shared resource is a shared memory resource.

3. The heterogeneous three-dimensional circuit stack as in claim 2, wherein the shared memory resource includes dynamic random access memory.

4. The heterogeneous three-dimensional circuit stack as in claim 2, wherein the shared memory resource includes non-volatile memory.

5. The heterogeneous three-dimensional circuit stack as in any one of claims 1 to 4, wherein at least one of the first processor or the second processor is an accelerator.

6. The heterogeneous three-dimensional circuit stack as in any one of claims 1 to 5, additionally comprising a third processor, wherein the third processor includes a graphics processor.

7. The heterogeneous three-dimensional circuit stack as in claim 6, wherein the third processor includes an accelerator and a graphics processor.

8. The heterogeneous three-dimensional circuit stack as in claim 7, wherein the interconnect logic enables coherent access to the shared resource for execution threads of the graphics processor.

9. The heterogeneous three-dimensional circuit stack as in claim 8, wherein the interconnect logic enables coherent access via programmatically assigned traffic classifications.

10. A system comprising:

a heterogeneous three-dimensional circuit stack including a first processor communicatively coupled with a second processor through one

or more through silicon vias, wherein the first processor and the second processor are graphics processors;

a first memory communicatively coupled with the first processor;

interconnect logic to communicatively couple the first processor and the second processor to a shared resource, the interconnect logic to couple with the first processor and the second processor via corresponding on-chip interconnects; and

wherein the interconnect logic includes bandwidth sharing logic to adjust bandwidth to the shared resource;

the shared resource includes a second memory to cache data to be received via the interconnect logic;

wherein the interconnect logic is to operate at a higher frequency than one of the first processor or the second processor; and

wherein the interconnect logic is to divide a message transmitted from the first processor or the second processor to another processor into a number of separate messages based on a frequency ratio of the interconnect logic, the first processor, the second processor and the another processor.

11. The system as in claim 10, wherein the shared resource is a shared memory resource including dynamic random access memory.

12. The system as in claim 10 or 11, wherein the first memory includes dynamic random access memory.

13. A method of interconnecting a heterogeneous three-dimensional circuit stack, the method comprising:

communicatively coupling a first processor to a second processor through one or more through silicon vias, wherein the first processor and the second processor are graphics processors, the first processor and the second processor communicatively coupled to a shared resource via interconnect logic via corresponding on-chip interconnects;

configuring bandwidth sharing logic to adjust bandwidth to the shared resource; and

configuring memory to cache data received via the interconnect logic,

wherein the interconnect logic is to operate at a higher frequency than one of the first processor or the second processor, and

wherein the interconnect logic is to divide a message transmitted from the first processor or the second processor to another processor into a number of separate messages based on a frequency ratio of the interconnect logic, the first processor, the second processor and the another processor.

14. The method as in claim 13, additionally comprising communicatively coupling a third processor to the interconnect logic, the third processor including an accelerator or a graphics processor.

15. An apparatus comprising:

means for communicatively coupling a first processor to a second processor through one or more through silicon vias, wherein the first processor and the second processor are graphics processors, the first processor and the second processor communicatively coupled to a shared resource via interconnect logic via corresponding on-chip interconnects;

means for configuring bandwidth sharing logic to adjust bandwidth

to the shared resource; and

means for configuring memory to cache data received via the interconnect logic,

wherein the interconnect logic is to operate at a higher frequency than one of the first processor or the second processor, and

wherein the interconnect logic is to divide a message transmitted from the first processor or the second processor to another processor into a number of separate messages based on a frequency ratio of the interconnect logic, the first processor, the second processor and the another processor.

16. The apparatus as in claim 15, additionally comprising means for communicatively coupling a third processor to the interconnect logic, the third processor including an accelerator or a graphics processor.