Benchmarking the Memory Interface Controller bus of the Cell processor

Nathalie Casati

EPFL

December 18, 2007

The Elements

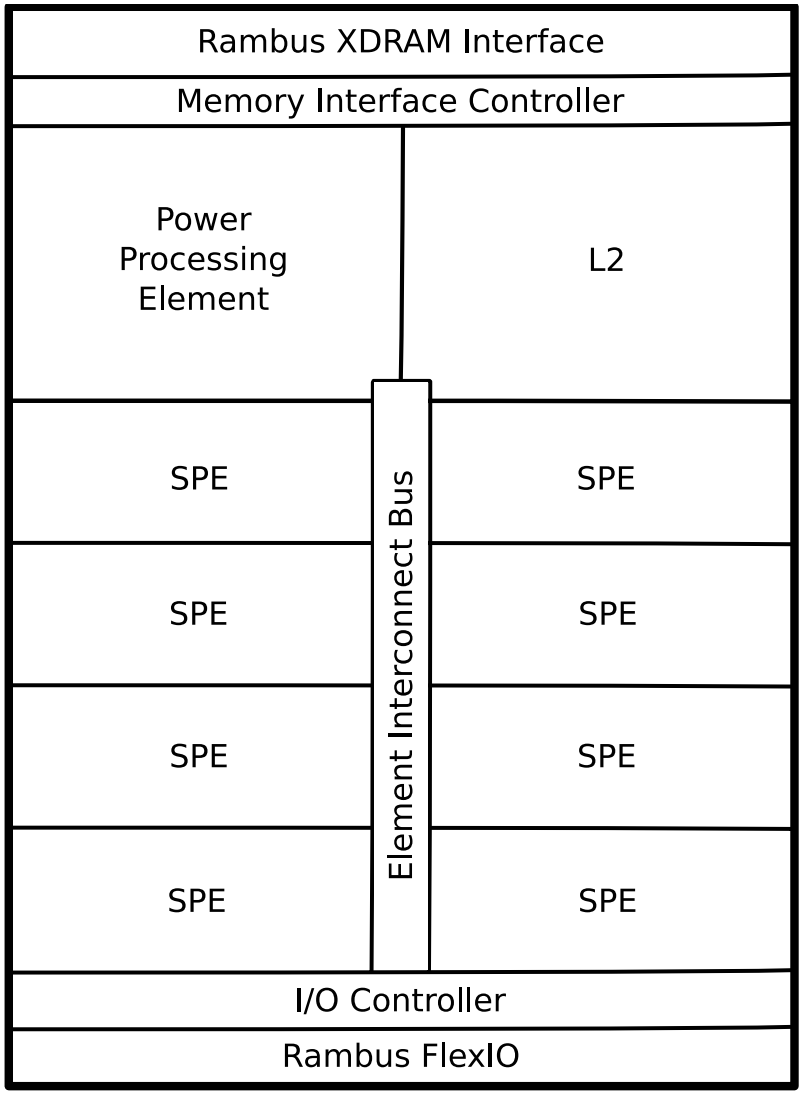

The main elements of the processor are :

- The Power Processing Element (PPE)

The Elements

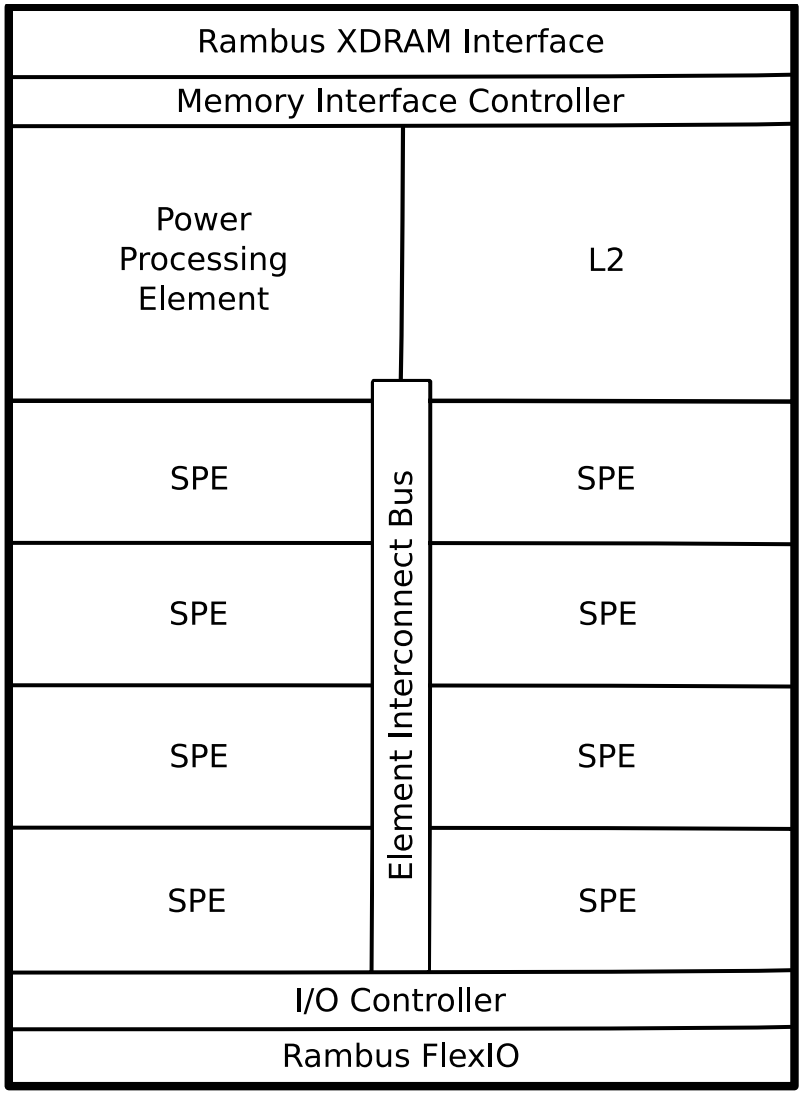

The main elements of the processor are :

- The Power Processing Element (PPE)
- The 6 (usable) Synergistic Processing Elements (SPEs)

The Elements

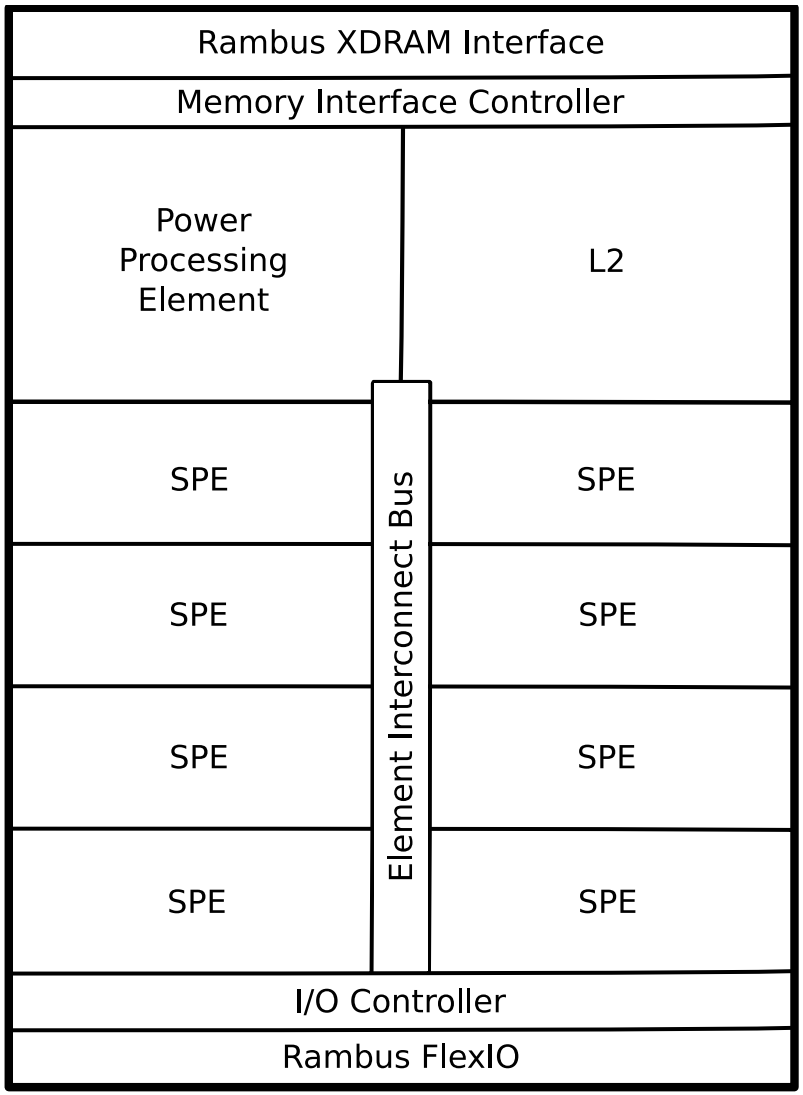

The main elements of the processor are :

- The Power Processing Element (PPE)
- The 6 (usable) Synergistic Processing Elements (SPEs)
- The Element Interconnect Bus (EIB)

The Elements

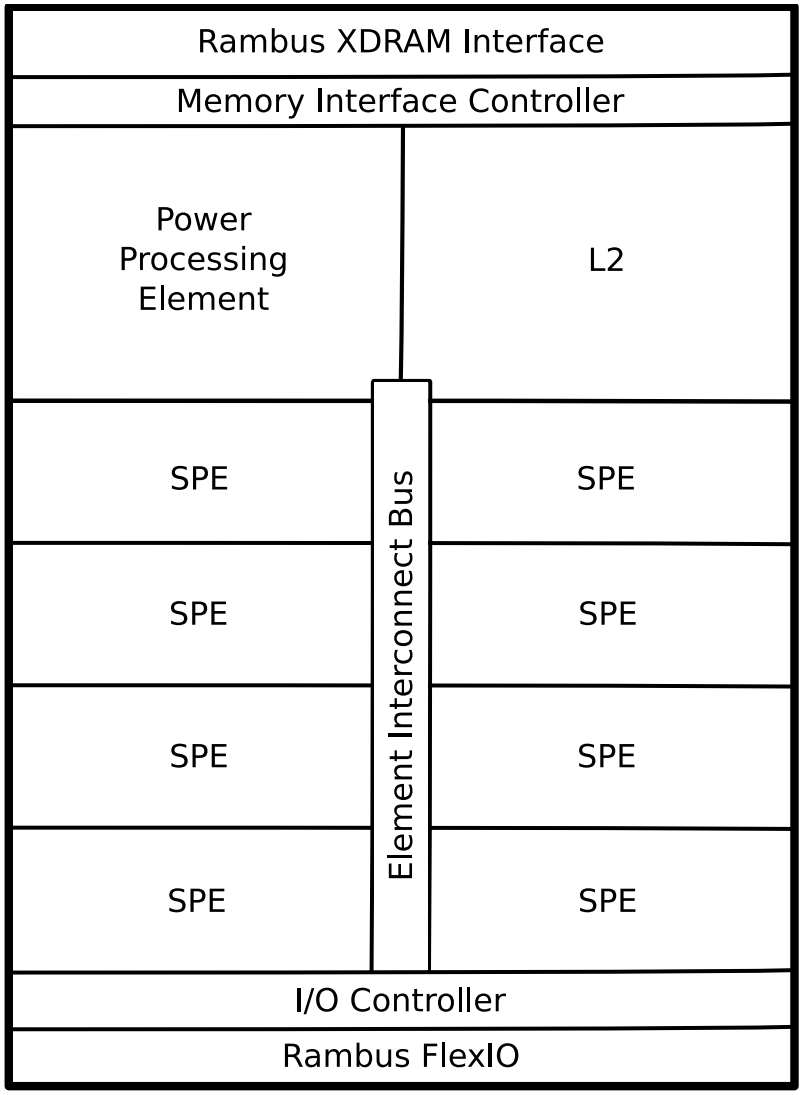

The main elements of the processor are :

- The Power Processing Element (PPE)
- The 6 (usable) Synergistic Processing Elements (SPEs)
- The Element Interconnect Bus (EIB)
- The I/O Controller (XIO)

The Elements

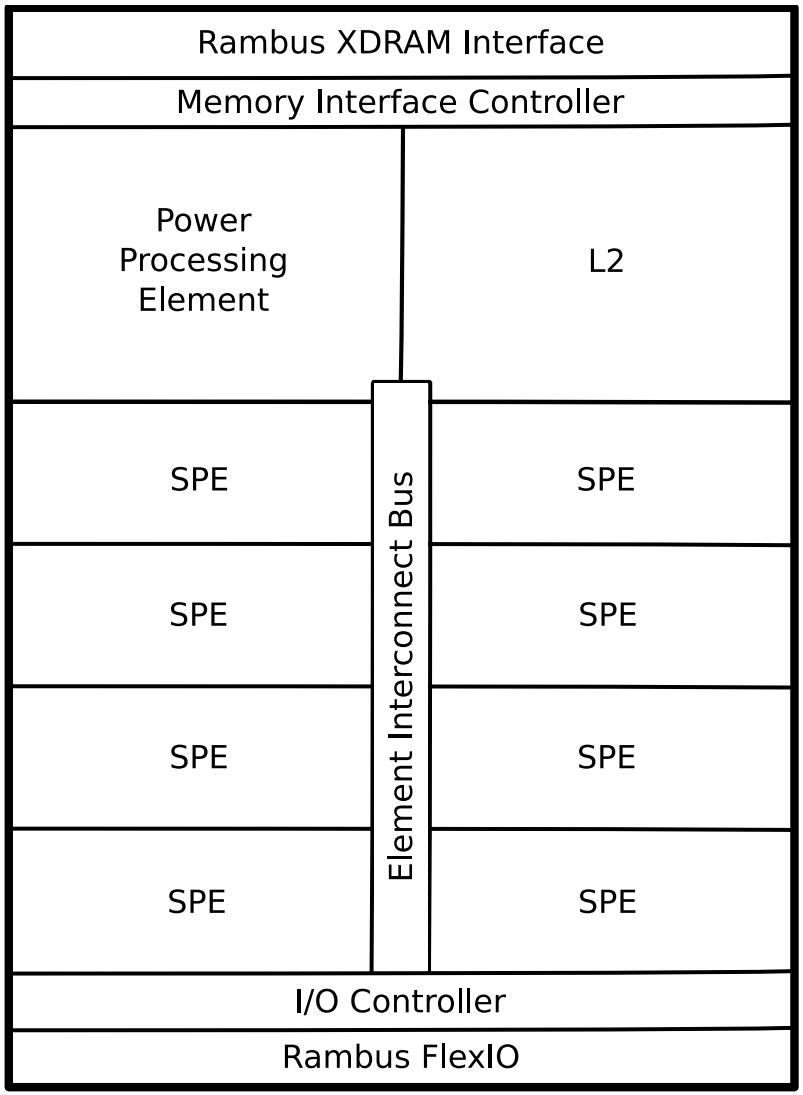

The main elements of the processor are :

- The Power Processing Element (PPE)
- The 6 (usable) Synergistic Processing Elements (SPEs)
- The Element Interconnect Bus (EIB)
- The I/O Controller (XIO)
- The Memory Interface Controller (MIC)

Other important facts

- Each SPE has a 256KB local store and no cache

Other important facts

- Each SPE has a 256KB local store and no cache
- Each element is connected to the EIB with 25.6 GB/s bandwidth (ingoing / outgoing)

Other important facts

- Each SPE has a 256KB local store and no cache
- Each element is connected to the EIB with 25.6 GB/s bandwidth (ingoing / outgoing)
- The 256MB main memory is connected to an external two channel Rambus XDR

Other important facts

- Each SPE has a 256KB local store and no cache
- Each element is connected to the EIB with 25.6 GB/s bandwidth (ingoing / outgoing)
- The 256MB main memory is connected to an external two channel Rambus XDR
- Each data transfer between SPEs and main memory is an explicit DMA operation up to 16KB

The main question

What happens if we use n SPEs at the same time while the MIC has *only* 25.6 GB/s bandwidth ?

(for a DMA get operation)

How to get the full bandwidth ?

- Sequential accesses read or write equal amounts of data to all memory banks

How to get the full bandwidth ?

- Sequential accesses read or write equal amounts of data to all memory banks
- Both effective address and the local storage address are 128-byte aligned

How to get the full bandwidth ?

- Sequential accesses read or write equal amounts of data to all memory banks
- Both effective address and the local storage address are 128-byte aligned
- Other factors like avoiding TLB misses

Bandwidth graphs

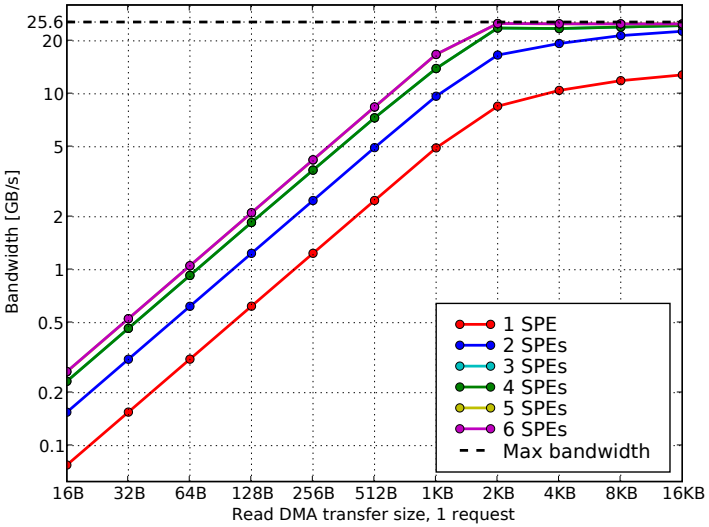

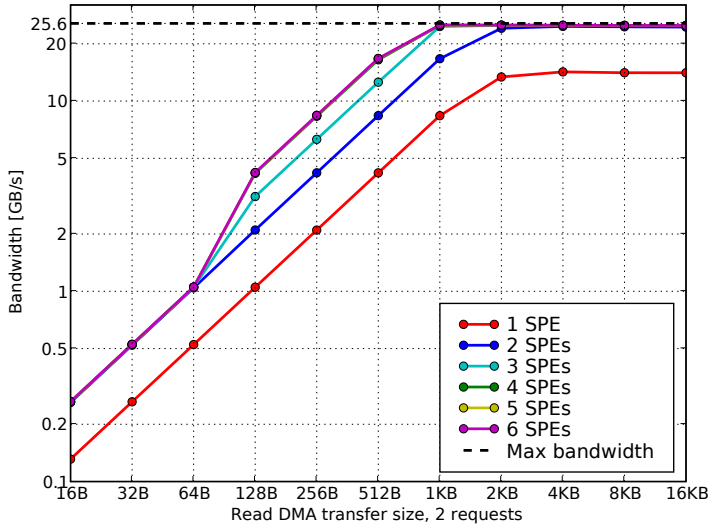

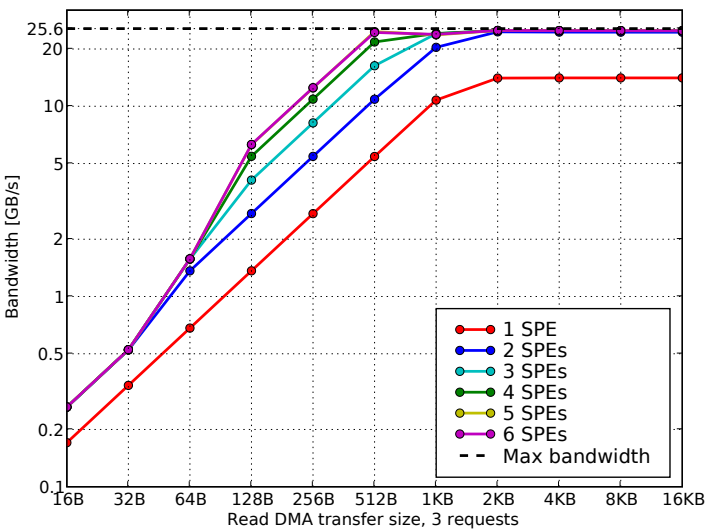

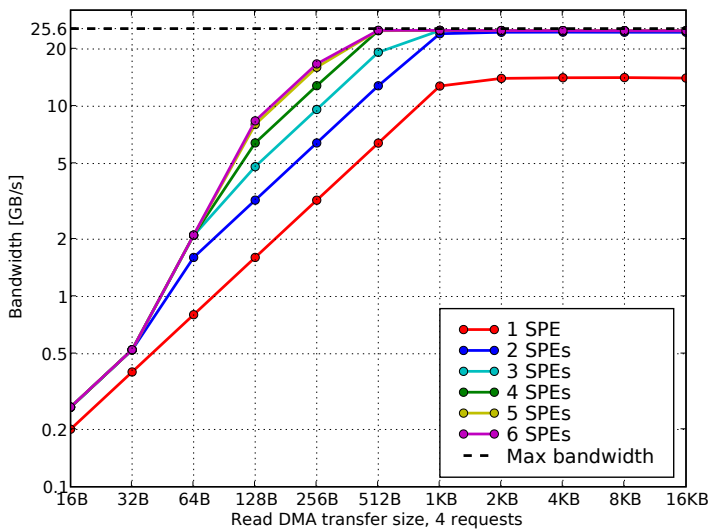

Bandwidth graphs

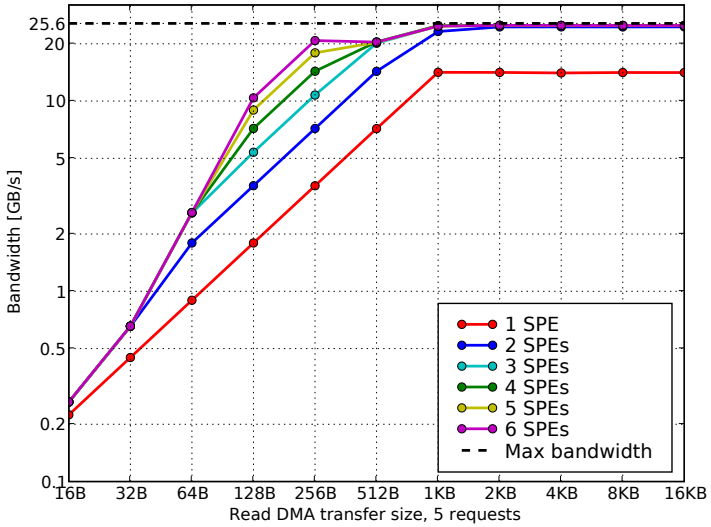

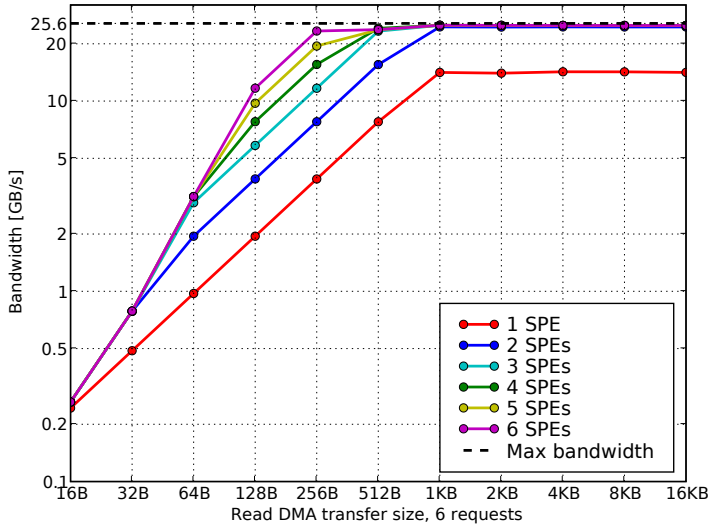

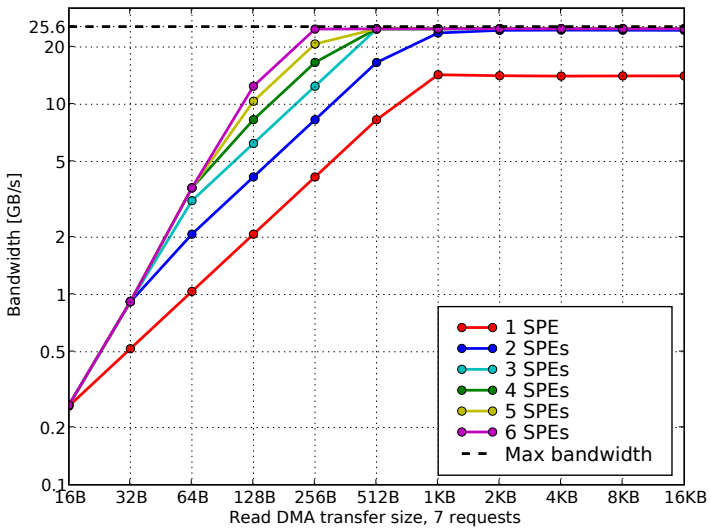

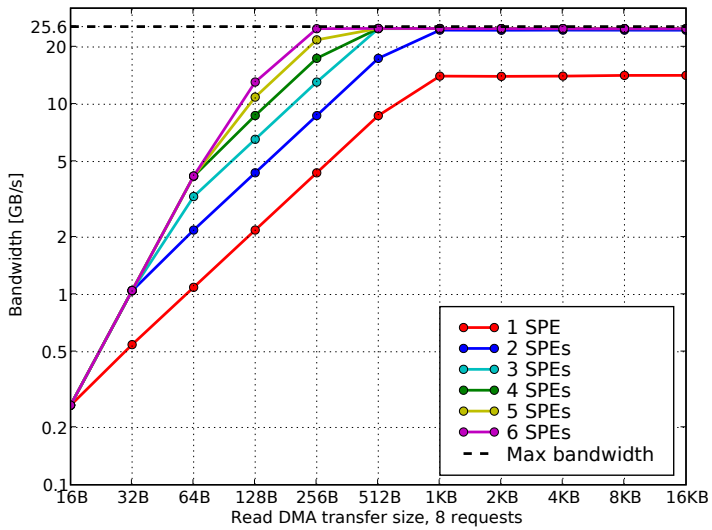

Results analysis and conclusion

- One SPE never gets the whole bandwidth (only about half)
  $\longrightarrow$ favours parallel accesses

Results analysis and conclusion

- One SPE never gets the whole bandwidth (only about half) $\longrightarrow$ favours parallel accesses
- The EIB is optimized for larger transfers

Results analysis and conclusion

- One SPE never gets the whole bandwidth (only about half) $\longrightarrow$ favours parallel accesses
- The EIB is optimized for larger transfers
- It is a good idea to use multibuffering

What's missing?

The same benchmark with a DMA put operation