P-16

Efficient On-Chip Networks for High-Performance Processors

Duraid MADINA ／ 理化学研究所基幹研究所

Out with the old...

- Supercomputers used to be built with “cut and paste”:
  - First, build a processor **chip**
  - Next, build a network **chip**
  - Stick them together,
  - Cut and paste, and you're done!
- Cut and paste was vital, because engineer time was scarce
- Fortunately, it never really mattered how good your network chip was, or how you cut and paste:
  - Meshes (maybe toroidal), n-cubes, different kinds of trees, crossbars, Kautz graphs, and various mixtures thereof: which you picked had only secondary effects because **data transport** was:
    - Fast, relative to processor speed
    - Used little energy, relative to processing
    - Used little space, relative to processors
  - So long as you didn't do a *terrible* job on your network chip and processor/network integration, application performance was always dominated by processing costs (time and/or energy)
    - And half your users would run embarrassingly parallel code anyway, leaving the network untouched
- Lots of great supercomputers built in this way:
  - Paragon, Exemplar, CM-5, SX-5, Meiko CS, T3E
    - Anyone who has spent time with two or more of these may enjoy reminiscing about how similar they really were, despite their *vastly* different networks
  - Supercomputers are still being built this way, **but**:

...in with the new!

- Future supercomputers will be built with a new “cut and paste”:
  - First, build a processor **core**
  - Next, build a network **core**
  - Stick them together,
  - Cut and paste, and you're done!
- Cut and paste still vital, because engineer time is still scarce
- Unfortunately, it **will** matter how you good your network core is, and just how you integrate them with the processor cores and with each other, because **data transport** is now:
  - Slow, relative to processor speed
  - Uses lots of energy, relative to processing
  - Uses lots of space, relative to processors

...and the relative cost of data transport is only going to increase

- The great supercomputers of tomorrow will have **tightly integrated processor and network elements**, because:
  - Pretending the data transport problem doesn't exist will hurt, unless you have embarrassingly parallel code
  - Running away from the data transport problem (e.g. through massive multithreading/other forms of “stream computing”) will hurt, unless you have embarrassingly parallel or embarrassingly linear code
- At RIKEN, we are developing **efficient on-chip networks**, a key technology to tackle the problem of data transport in future supercomputers