~~System Level~~ Interconnect Prediction

Then, now, and future

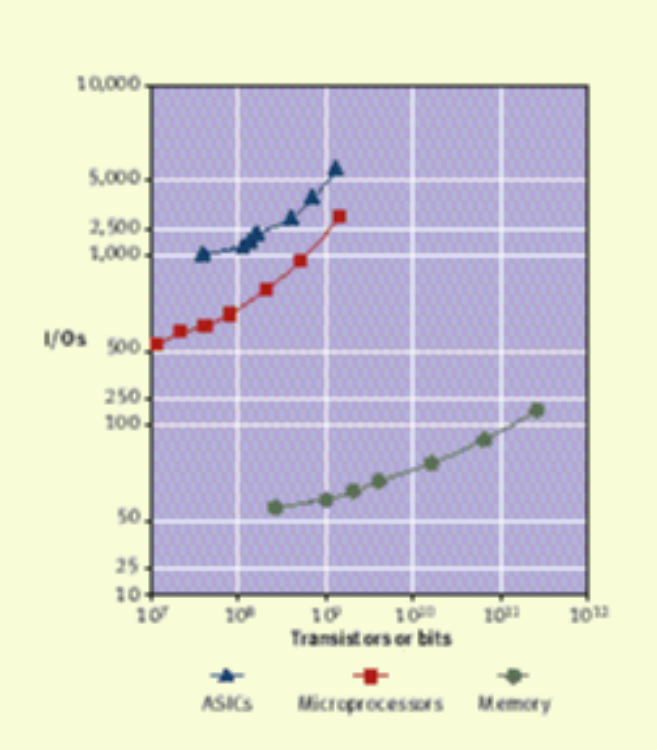

Interconnect estimates work for systems that look this:

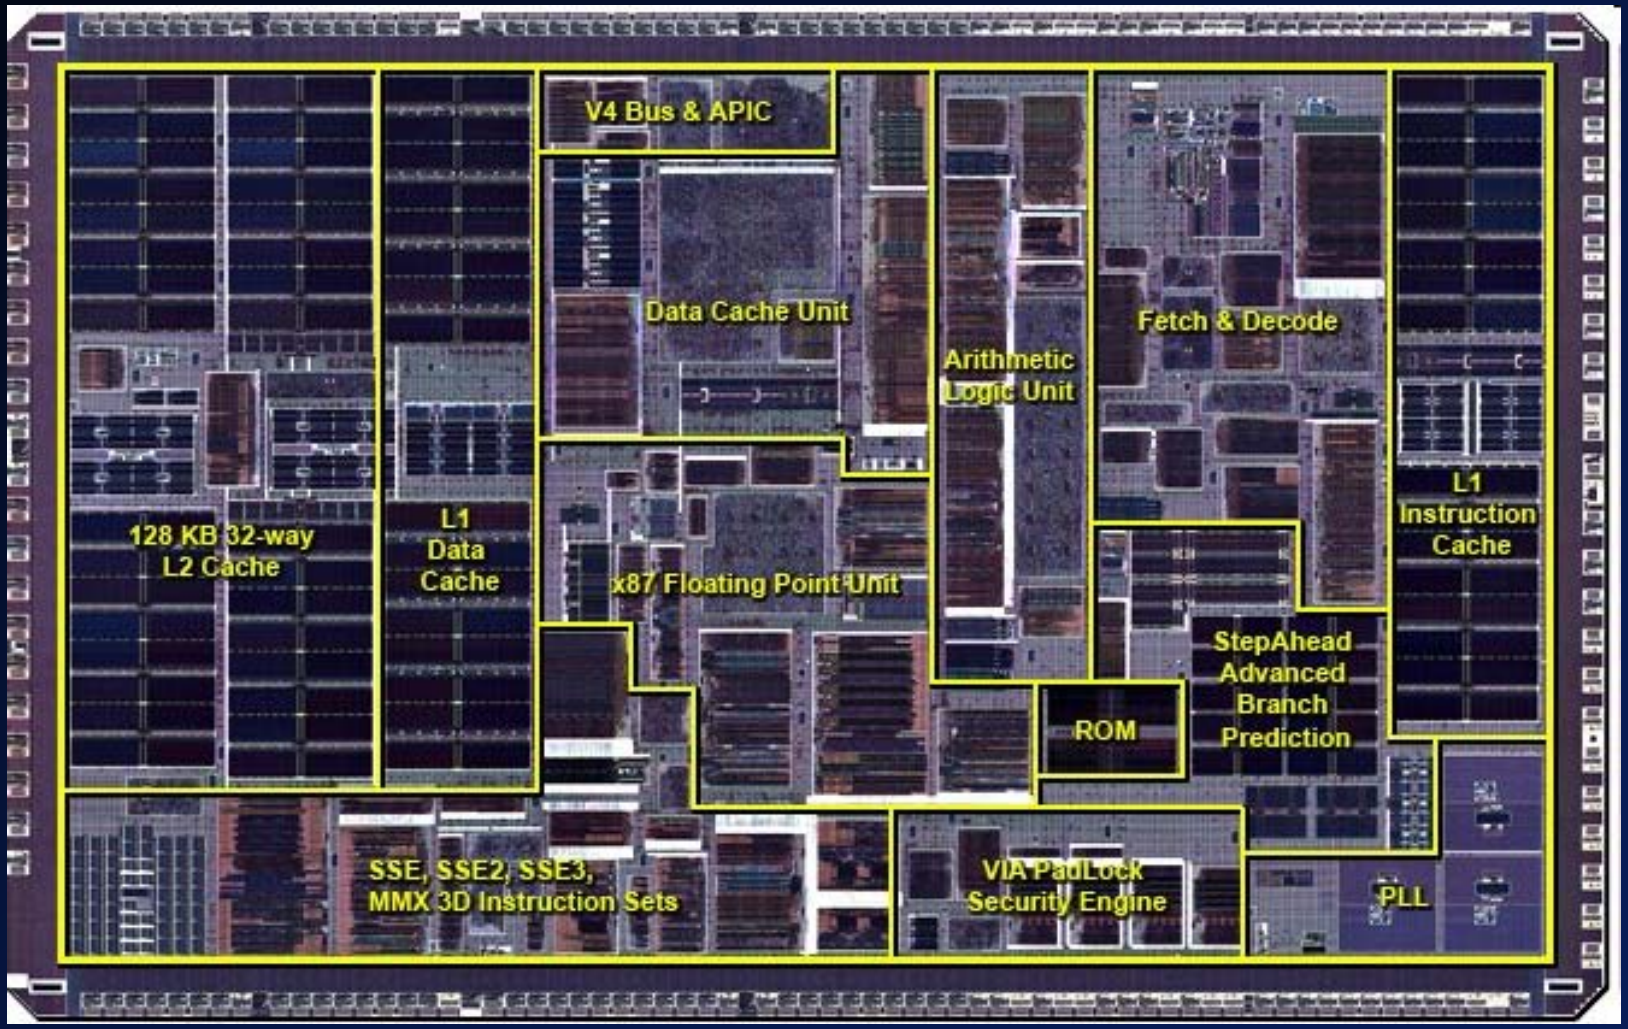

C7 die layout from linuxdevices.com

Interconnect analysis started with Rent's rule

- Not specifically system level, though it works for some systems
  - In particular, works well for CPUs
- Generalized to wire length properties
- Work well for pure scaling
- But, has a number of limitations leading to less relevence

Even on a single chip, this is breaking down

- Fast connection costs ~8x a slow one
- Exact same design may have very different interconnect needed depending on performance

Name ‘SLIP’ implies we can do this, but we can’ t

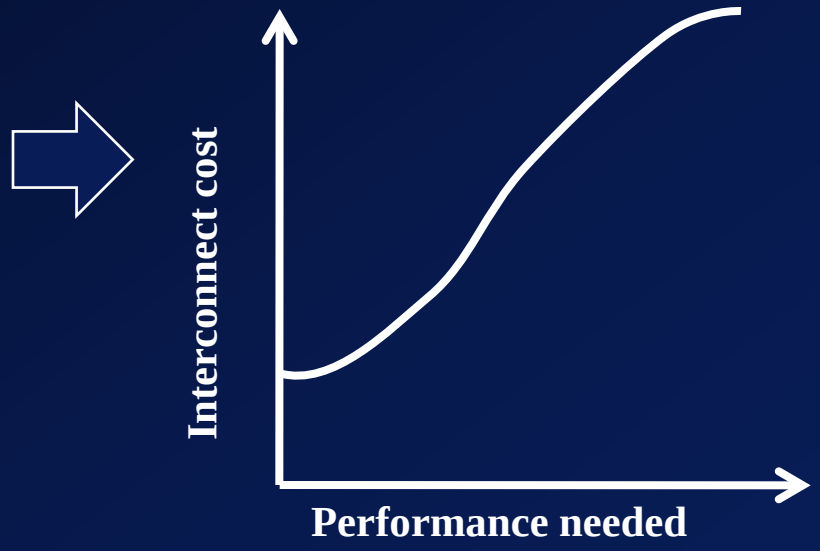

But a system is heterogeneous, and we have even fewer interconnect tools

- Wide variety of scales and needs of interconnect

Interconnect in a wide variety in systems

- Small, fast, single ended on chips
- Bigger on PC boards
- Backplane may be optical fiber
- Long runs - twisted pair, coax, optical, RF, networks, etc.
- Each has own expense and technical issues
- Commercial and political considerations

This is where SLIP should go to earn its name

- Very few tools in this area now
  - Done by system architects, using previous experience
- Functions that each block performs depends on what interconnect is chosen
- Estimating costs, predicting performance, finding good configurations all are difficult and time consuming, since both interconnect and blocks change
- Big problem and big opportunity

Papers for next 10 years

- “A language for specifying interconnect options”
- “Combining interconnect prediction and logic synthesis to minimize system cost”
- “Optimizing connection costs in mixed electronic/biological systems”