Schedule

- **2:00-3:05 (Diana Marculescu)**
  - Trends and issues in clock distribution
  - Motivation for GALS, synchronization issues, deadlock prevention, possible inter-clocking domain communication schemes
  - GALS processors: power/performance evaluation
- **3:05-3:30 (Dave Albonesi)**
  - GALS processors: power/performance evaluation (cont'd)
- **3:30-4:00 Break**
- **4:00-4:40 (Dave Albonesi)**
  - Workload characterization and impact on the use of fine grain speed/voltage scaling
- **4:40-5:45 (Alper Buyuktosunoglu, Pradip Bose)**
  - Case study - LPX, an IPCMOS based processor
- **5:45-6:00 (Diana Marculescu)**
  - Looking in the crystal ball: Where will partial asynchrony be used?
  - Concluding remarks

Where will partial asynchrony be useful?

Let's remember the motivation…

- **Current design roadblocks:**
  - Increasing power density
  - Increasing design complexity
- **Shortening time-to-market require:**
  - Design and verification tools and methodology
  - Solutions for design testability and reliability
- **Drivers for minimally clocking:**
  - Design reuse (IPs independently designed)
  - Minimize design effort for global clock distribution
  - Allow fine-grain application-driven adaptability (voltage and speed scaling)

Ahead: possible partially asynchronous systems

- **Clustered architectures**
  - Take advantage of localized computation
  - Presumably less communication overhead
  - Decoupling will enable fine-grain voltage/speed adaptation
  - Possibly in conjunction with dynamic instruction steering based on criticality
  - Regularity may help with system variation, and hence clock skew and local clock speeds

Ahead: possible partially asynchronous systems

- **Clustered architectures**
- **General multiprocessor-based systems (CMPs)**
  - Are naturally decoupled
  - Regularity should help in this case too…

Ahead: possible partially asynchronous systems

- **Clustered architectures**
- **General multiprocessor-based systems (CMPs)**
- **Heterogeneous Systems-on-Chip**
  - Tile-based architecture may improve interconnect problems
  - Separation of concerns: Computation vs. Communication
  - Important: design of Networks-on-Chip (NOCs) among tiles running at different local clocks
  - Again, workload driven adaptation may be important

Ahead: possible partially asynchronous systems

- **Clustered architectures**
- **General multiprocessor-based Systems-on-Chip (CMPs)**
- **Heterogeneous Systems-on-Chip**
- **In the long run, maybe all systems will have to be partially asynchronous**
  - Process and system variation may reach or exceed 100%
  - Non-determinism and asynchrony will become prevalent:
    - Instead of fixed signal arrival times…
    - …You will have deal with "There's a chance that the signal will arrive during this interval!"
  - ➔ Drastic changes in the design process!

Conclusion

- **Increasing complexity, design effort and process technology adverse effects on global clock distribution will necessitate alternative design styles**
- **Minimal clocking offers lower power consumption, but performance may be affected**
- **In conjunction with fine-grain DVS, minimally-clocked machines may be more power efficient than their fully synchronous counterparts**
- **Not yet fully explored:**
  - Impact of global or local asynchrony on other types of architectures (e.g., clustered)
  - Effect of smaller clocking area on clock skew
    - Less process variation, hence less clock skew
    - Will regularity help? (less system variation)

Thank you!

More information:

http://www.ece.cmu.edu/~enyac

http://www.ece.rochester.edu/~albonesi/acal

http://www.research.ibm.com/lowpower/

dianam@ece.cmu.edu, albonesi@ece.rochester.edu, alper@ece.rochester.edu, pbose@us.ibm.com