Computer Architecture Reading Group Notes

Date: 2/19/04
Discussion Leader: Vicky
Notes: Ben
Topic: Biocomputing
Papers:

1. Seth Copen Goldstein and Mihai Budiu. NanoFabrics: Spatial Computing Using Molecular Electronics. In *Proc. 28th Annual International Symposium on Computer Architecture*, pp. 178-191, June 2001.
2. M. Ogihara, A. Ray, and K. Smith, DNA computation--a shape of computing to come, *SIGACT News* 28(3), pp. 2-11, 1997.
3. M. Ogihara and A. Ray, Biomolecular computing--recent theoretical and experimental advances, *SIGACT News* , 30(2):22-30, 1999.

Spatial Computing Using Molecular Electronics

Summary:
CMOS can't scale forever. This paper proposes using nanoscale technology including wires and diodes to create logic networks. In this technology, special consideration is required to cope with defects, which are expected to be commonplace. The authors propose regular structures as the basis for a reconfigurable computing fabric. There are some simplistic simulations to demonstrate plausibility. While the work is clearly preliminary and does not address all problems, the paper presents an interesting solution to anticipated issues in future technology.

Discussion:

- Why does reconfigurable computing apply here? Due to defect distributions.
- Note that they don't use parallelism available in reconfigurable computing.
- They neglect global interconnect/routing issues. Routing distance relative to device dimensions is expected to grow, so this issue is exacerbated.
- Is it realistic to assume that 3-terminal devices wouldn't be available by the time the technology becomes viable?
- Good power consumption ~2.5W claimed.
- Vicky mentions other studies suggesting this type of technology has been proposed for use as memory only, since memory makes up an increasing portion of processing cores.
- Clock frequencies reported weren't so impressive for a theoretical future technology (100MHz – 1GHz).
- How would configuration/encoding work? They mention parallel configuration but they don't explain the complications very well.
- 750k clusters per cm^2, cluster = 128 blocks. Better than FPGA.

DNA Computation – A Shape of Computing to Come
and
Biomolecular Computing – Recent Theoretical and Experimental Advances

Summary:
These papers provide a background on DNA manipulation techniques that can implement computation, particularly as a vehicle for brute-force solution of NP problems. The basic approach is to create an initial "stew" of genetically-encoded possible solutions and then apply genetic operators to eliminate the incorrect answers. This provides a highly parallel form of computation. The major limitations to this approach are error rates, processing time, and the scaling of the volume of DNA needed.

Discussion:

- Rates of reactions make execution take days. Not very automated just yet.
- How do you program it? Seems to require algorithms to follow a fairly rigid structure without temporary variables.
- What are the killer apps? Highly parallel, error tolerant – things like code cracking.
- Exponential time in conventional algorithms becomes exponential space in DNA. The volume constraint becomes prohibitive at a point at which benefit over conventional computing is still unclear.
- They don't have any I/O.
- Instructions are encoded using enzymes, inhibitors.
- Confusion about how parallelism is expressed here.
- Error computation based on copy errors, but what if it isn't in original soup? Requires huge "soup" to reduce chance of missing solution. Initial soup is expected to have redundancies.