Bending Light for Multi-Chip Virtual PRAMs?

Uzi Vishkin

Igor Smolyaninov

Both: ECE. UMIACS. Maryland

Optics Group (MOG)

A SINGLE-LAYER WAVEGUIDE MODEL

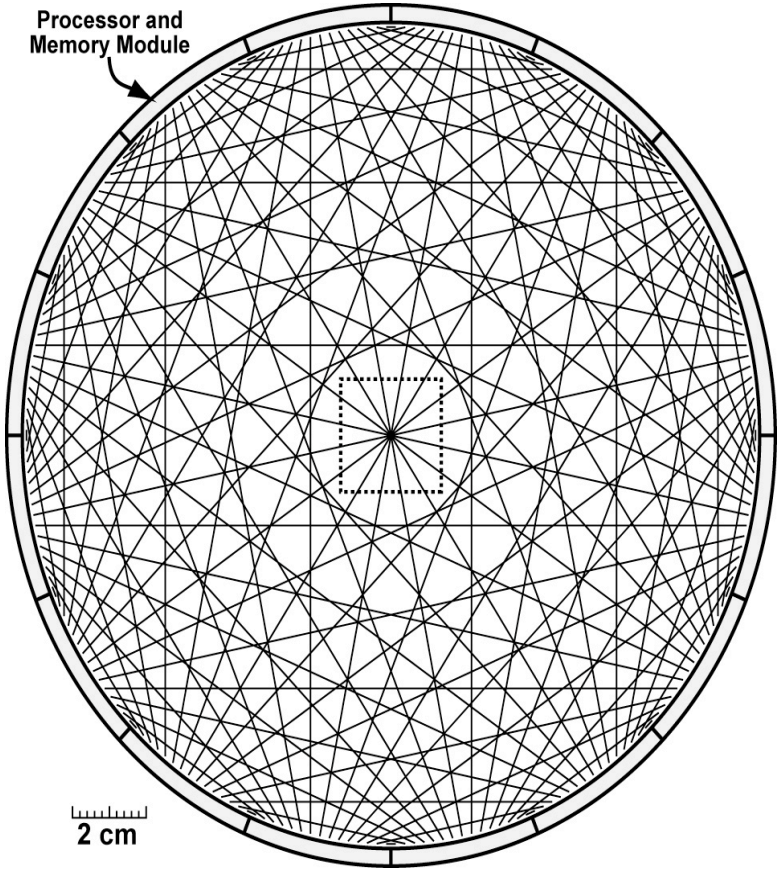

ALL-TO-ALL where

- bending optical comm channel: limited
- 2 OPCs cross? 90 degrees or close
- <=2 OPCs can cross at the same point
- distance between any two crossing points: not too small
- distance between 2 OPCs not too small (unless crossing)

Assume: 20cm diameter

TECHNIQUE to fit model

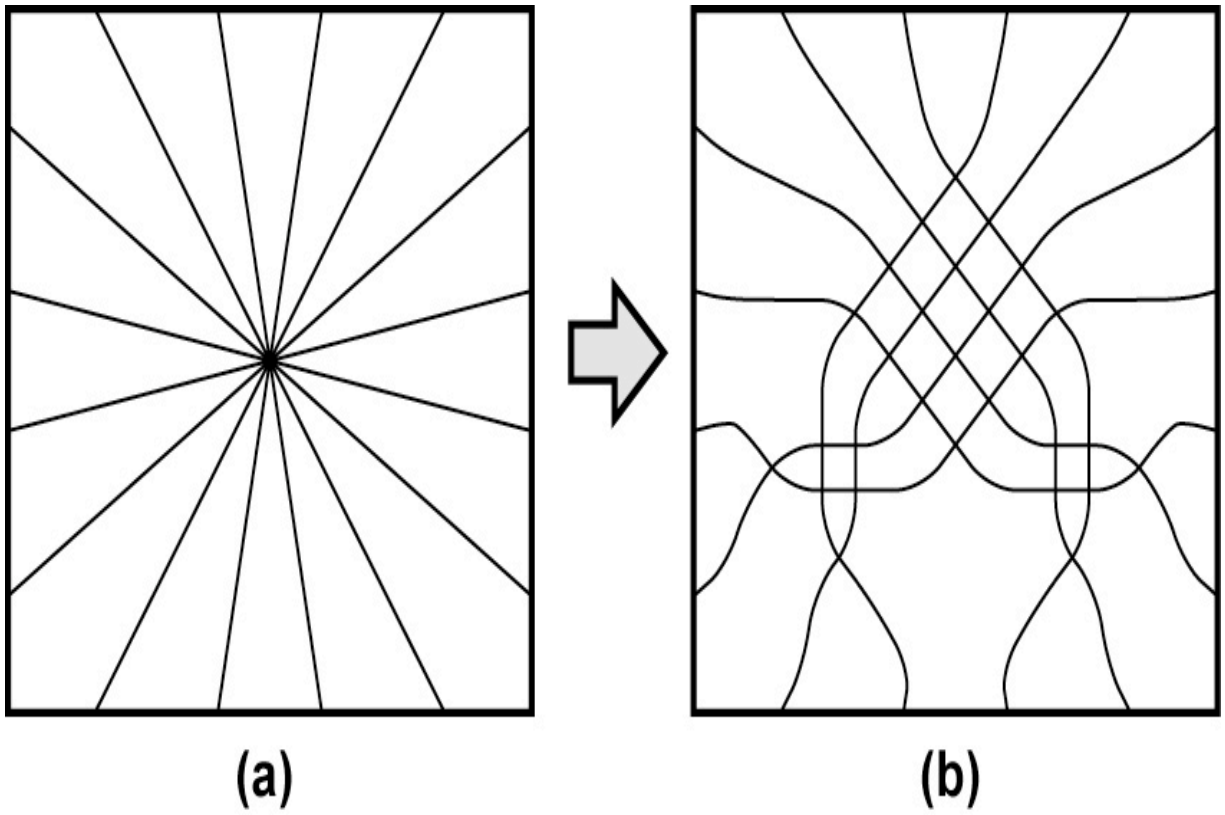

Technology Discussion 1

- Optics: massive, longer data movement. Electronics: everything else.
- Data rates: complex. Indirectly addressed.
- Conversion time: will be dominated by flight time. Already under .025ns/bit.
- Thermo-modeling (-signal driving, +conversion)
- Diameter? 30cm/ns speed of light.
- Which waveguide technology: fabrication, performance. Need to mature towards the current application domain. E.g.: emitters&detectors integration into silicon (separate GaAs plane?)
- How many crossings? Radiative/scattering loss? See paper.
- 2-layer implementation
- Inexpensive packaging?!

Architecture Discussion 1

- Focus on our motivation. There could (and should) be others.
- PRAM: abstract model for studying parallel algorithms. Assumes: many memory accesses to the shared memory can be satisfied within the same time as one. Huge algorithmic theory; second only to serial theory.
- Open question: can PRAM-like algorithm be helpful in practice.
- CS Mantra in 1990s: “PRAM is unrealistic”

Architecture Discussion 2

- Explanation only for multi-chip multi processing (note: this paper questions that…).
- UMD "PRAM On Chip" project: fit, say, 64 memory modules + 64 processor clusters on chip. Explicit multi-threading (XMT) architecture:
- Von-Neumann apparatus (stored program coupled with program counter) extended to accommodate parallel programming and **low-overhead** parallel execution
- HW implementation of Fetch-And-Add.

Architecture Discussion 3

- Single-program multiple-data (SPMD) programming of PRAM-like algorithms along with "independence of order semantics“.
- memory architecture: memory modules provide 1st level cache for a hashed partitioned shared memory. Note: cache coherence is defined away.
- Not enough time to present in sufficient detail, BUT the interesting question for this presentation is:

Architecture Discussion 4

Compare performance of:

- an all-electronic 90nm PRAM chip with 64(?) processor-cluster, (64?) memory modules and all-electronic interconnection network, against
- 64 chips in 130nm, each with one cluster and one module (small and cheap!), coupled with optical interconnection network

Technology Discussion 2

- Success story. Model of VLSI fabrication: possibly expensive template but cheap copies.
- Motivated the current approach, but can the vision be brought down to fabrication of the interconnection network and the overall packaging?

What Next (technology)

- Need to motivate optics to mature in our direction.
- Reach out to Optics community. Post NSC deadline news: Can Optical Interconnection Networks Lead to Cheaper High-Performance Multiprocessors? Accepted to Optics East, S.P.I.E. - The Intl Society for Optical Engineering, 10/2004.

What Next (architecture)

- Study