

# Pressures the Radiation Transport Problem Places on Future PIM- Based Supercomputer Designs

Thomas Christopher  
tc@tools-of-computing.com

Funded by Sandia National Labs  
Contract #263162

# Are PIMs Worth the Risk?

- Can PIMs run the important ASCII applications?
- How fast?
- How cost effectively?
- How will they in the future?

# Overview

- “Back of the envelope” calculations to determine the requirements certain ASCI applications impose on PIM designs.
  - Bits to store
  - Bits to move
  - Instructions to execute
- Project into future using ITRS tables.

# Applications To Consider

- Compare fixed cell size (e.g. CTH shock physics) vs. generation-dependent cell size (e.g. SWEEP3D radiation transport).
- Look more closely at radiation transport.

# Recoding the Algorithms

- CTH: already 3D decomposition.
- SWEEP3D:
  - 2D decomposition requires recoding as 3D.
  - Running within cache is a problem.

# Generations of CTH-like Problems

- Double the number of cells in each dimension.
- Cut the step time in two.
- Program to complete in same time.
- Space grows 8 times.
- Computation grows 16 times.
- Communication grows 8 times.


# Generations of Radiation Transport Problems

- Double the number of cells in each dimension.
- Double the number of angles.
- Double the number of energy levels.
- Cut the step time in two.
- Space grows 32 times.
- Computation grows 64 times.
- Communication grows 32 times.

# Projecting problem generations


- Look at fractional growth of new chip:
  - Memory, cycles, and off-chip bandwidth.
  - Fixed proportion of chip devoted to processors.
- Fractional growth in problem parameters:
  - Holding number of nodes constant, and
  - Assuming all chip memory is devoted to data, or
  - Assuming 50% of initial chip memory devoted to data, as is all the increasing memory.
- Can the program complete in the same time?

# Chip capacity



## Growth in processing per usable processor

Fraction growth




2001 2002 2003 2004 2005 2006 2007

Year

- ◆ Chip cycles per sec.
- CTH
- ▲ Rad. Trans.
- × CTH, 50% init .mem overhead
- \* Rad. Trans., 50% init. overhead

## Growth in processing per usable processor

Fraction growth



2001

2003

2005

2007

2013

Year

—◆— Chip cycles per sec.


—■— CTH

—▲— Rad. Trans.


—×— CTH, 50% init .mem overhead

—\*— Rad. Trans., 50% init. overhead

## Communication grows



# Communication grows



# A Huge Radiation Transport Problem

- Compare
  - System costs
  - Sweep times
  - Cost effectiveness
- For systems
  - MPU+DRAM
  - Pure PIM
  - PIM+DRAM

# Problem size

## Parameter

## Explanation

$N_a=625$

angles/octant

$N_s=2$

particle species

$N_e=1000$

energy levels

$D=1000$

dimension, number of  
cells= $D^3=1,000,000$

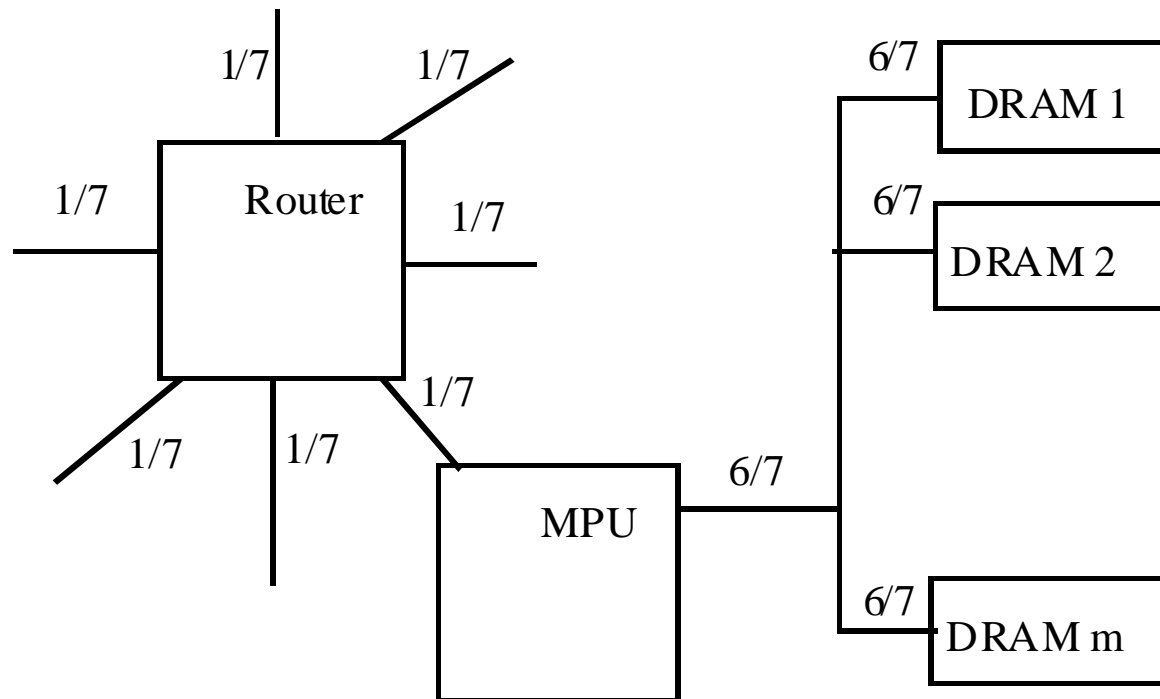
$N_{\text{sweep}}=1,250,000$

number of floats  
sent/sweep by a cell

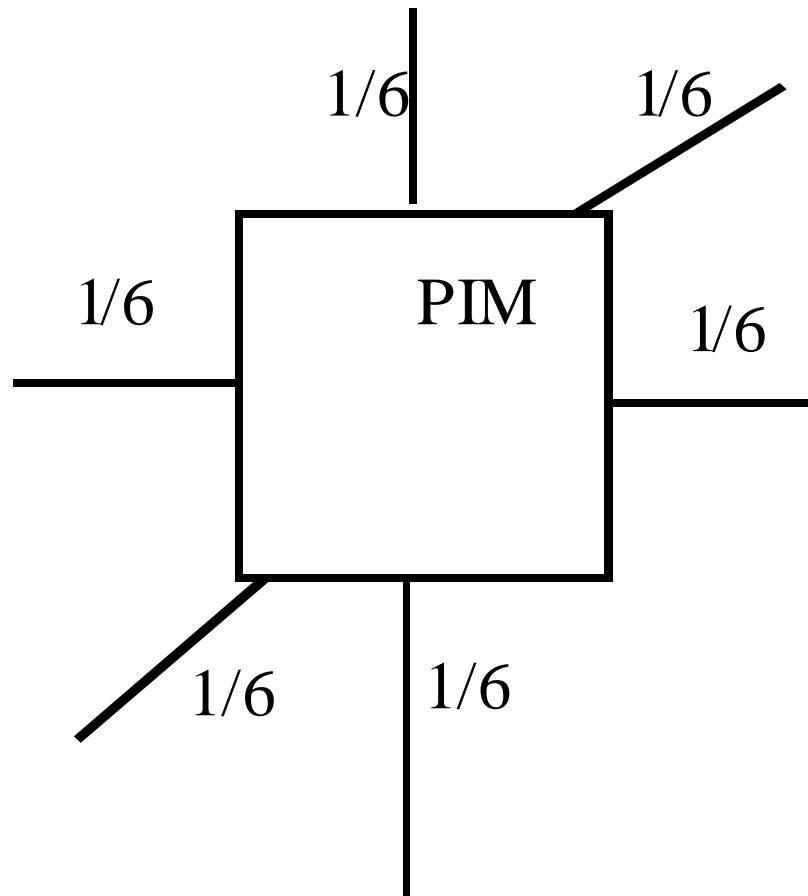
# How big is that?

- The overall space requirements for the  $10^9$  cells is  $1.6 \times 10^{17}$  bytes. At 4 gigabits per chip, that is 380,000,000 DRAM chips, containing 117,800 square meters of silicon, or a bit over **29 acres** or **22 football fields**.
- With DRAMs at \$5 apiece, the machine would cost \$1.9 billion in memory chips alone.

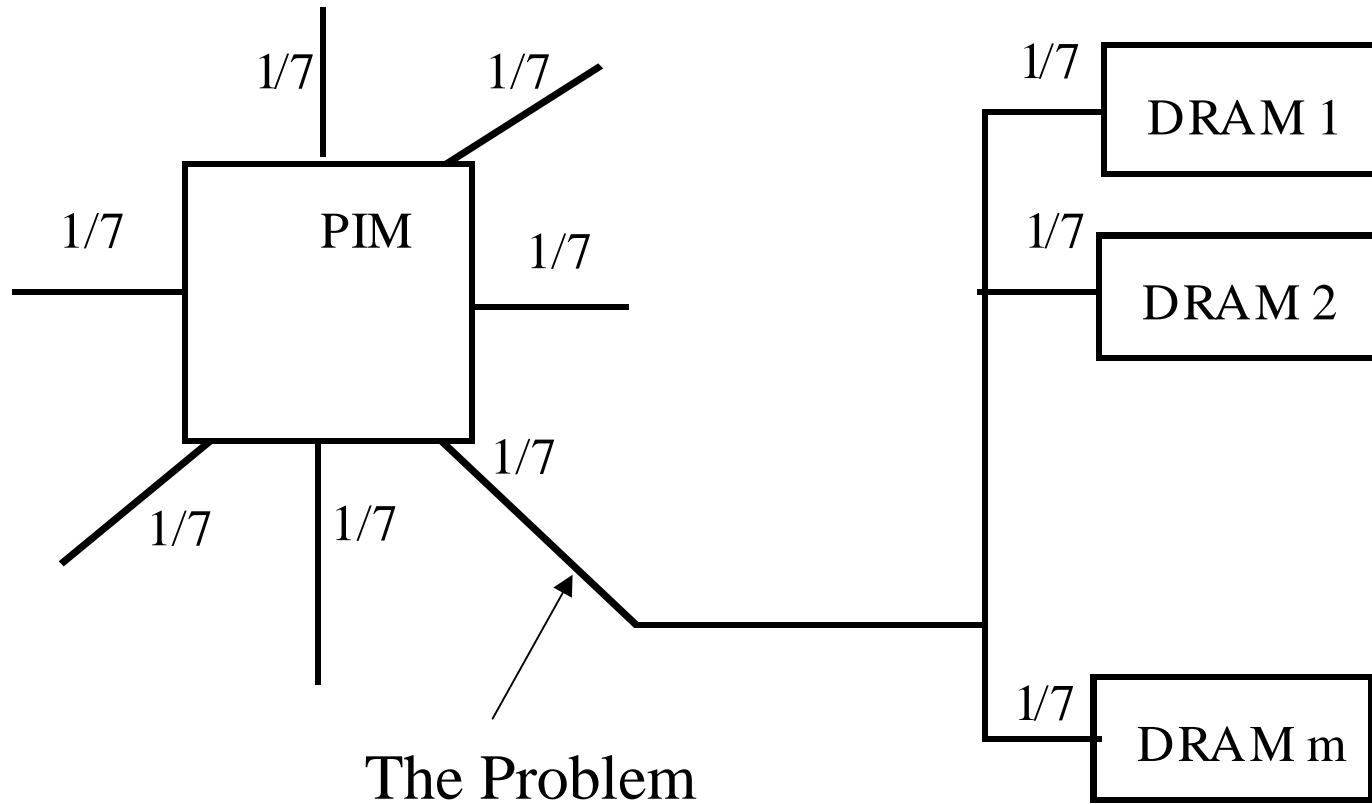
# “Red Storm”



# Pure PIM




# PIM + DRAM




# Prices assumed

- price per DRAM \$5.00
- price per MPU \$150.00
- price per router \$300.00
- price per PIM \$300.00


### Cost of minimal system




### Sweep time




### Cost per sweep



# System costs



## Costs per sweep



# Conclusions:

## Growth of problem sizes

- Near term, CTH should be able to grow with chip memory capacity without slowing down. Radiation transport may be reaching the limits of processor cycle speedup. Bandwidth may be okay.
- Longer term, radiation transport may strain processor cycles. Everything may be running out of bandwidth off-chip.

# Conclusions:

## Huge R. T. Problem

- Unless PIMs become commodities
  - Pure PIM system costs will be prohibitive.
  - Solution costs on PIM and conventional systems will be about the same.
- If PIMs become commodities
  - System costs become about the same.
  - PIM-only solution costs become an order of magnitude lower than MPU+DRAM.

END