



Evaluation of Active Graph Applications on the MTA-2

Shannon Kuntz, Jay Brockman,
Peter Kogge, Matthias Scheutz
University of Notre Dame

Mark James, Ed Upchurch
NASA JPL

John Feo
Cray Inc.


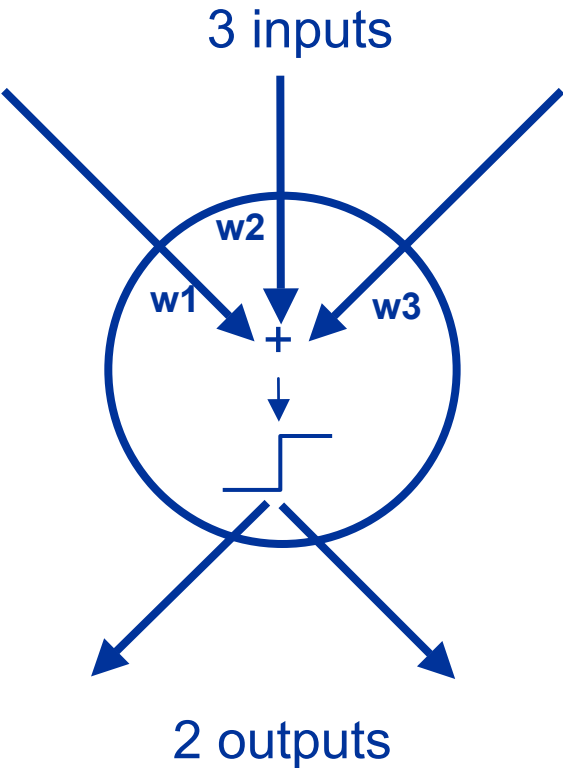
◆ Goals

- Examine massively-parallel lightweight multithreading for modeling systems of intercommunicating agents
- Look at tradeoffs between programming productivity and execution efficiency for active graphs on MTA-2

◆ Outline

- Definitions and example applications
- Representation of active graphs
- Experiments and results
- Conclusions and future work

- ◆ Definition
 - Each node (possibly edge) tied to a distinct thread
 - Flow of data through producer/consumer relationships between threads
- ◆ Thread state
 - input mailboxes
 - working registers
 - pointers to upstream/downstream nodes
- ◆ Key issues
 - exploiting parallelism
 - managing synchronization





Shine: Spacecraft Health Inference Engine

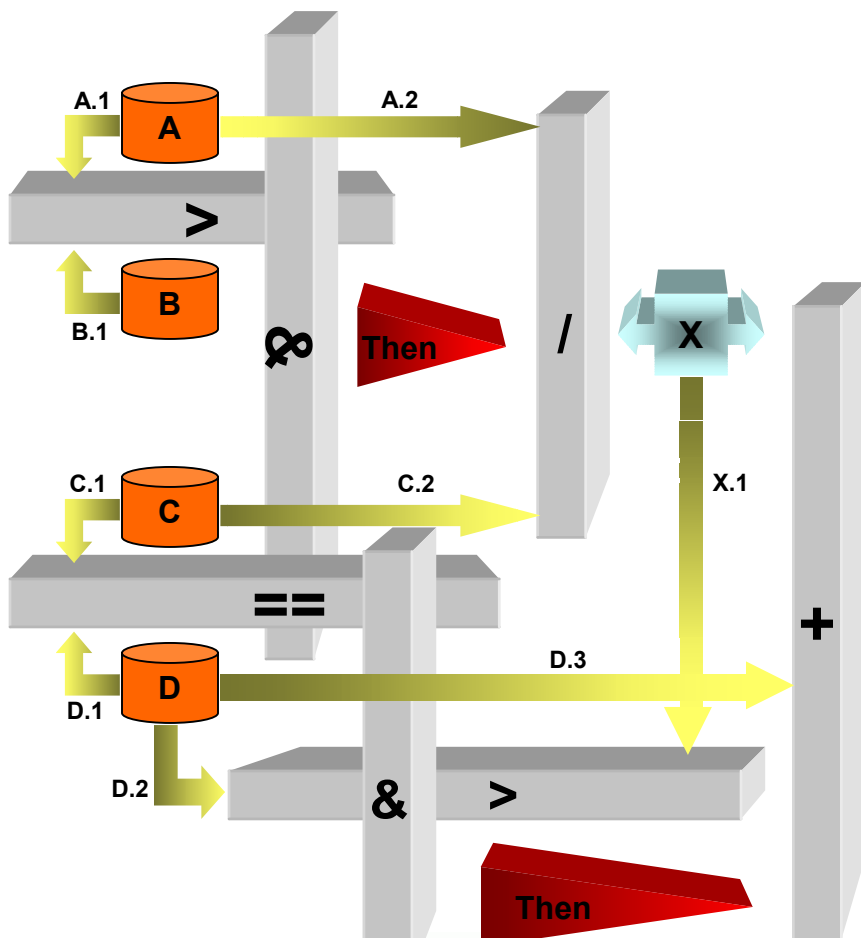


Rule 1:

If Sensor A > Sensor B & Sensor C == Sensor D
then Var X = Sensor A / Sensor C

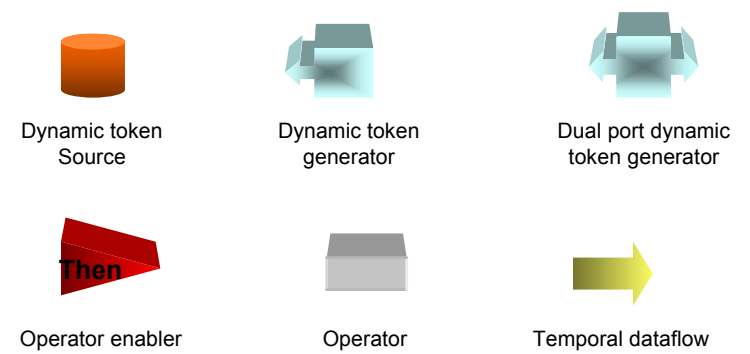
Rule 2:

If X > Sensor D & Sensor C == Sensor D
then Var Z = Var X + Sensor D



- All rules are analyzed respect to one another for all possible global interactions with maximal sharing.
- A sophisticated mathematical transformation, based on graph-theoretic data flow-analysis is introduced, that reduces the complexity of conflict-resolution during the match cycle from $O(n^2)$ to $O(n)$.
- The underlying structure is mapped into temporally invariant dataflow elements.
- The final representation is either executed in a debugging environment or translated to a variety of target languages such as ADA, C and C++.

Legend



Interface Information Procedures

Edit Delete abc Button abc Slider On Off Switch abc Chooser abc Monitor abc Plot abc Output abc def ghi jkl Text

NetLogo

pH level 7.4 = 70% Nicotine Ionization

Setup Go

thickness 25

Release Nicotine

number-of-pores 77

dosage 21 mg

number-of-reactants 101

Smoking None

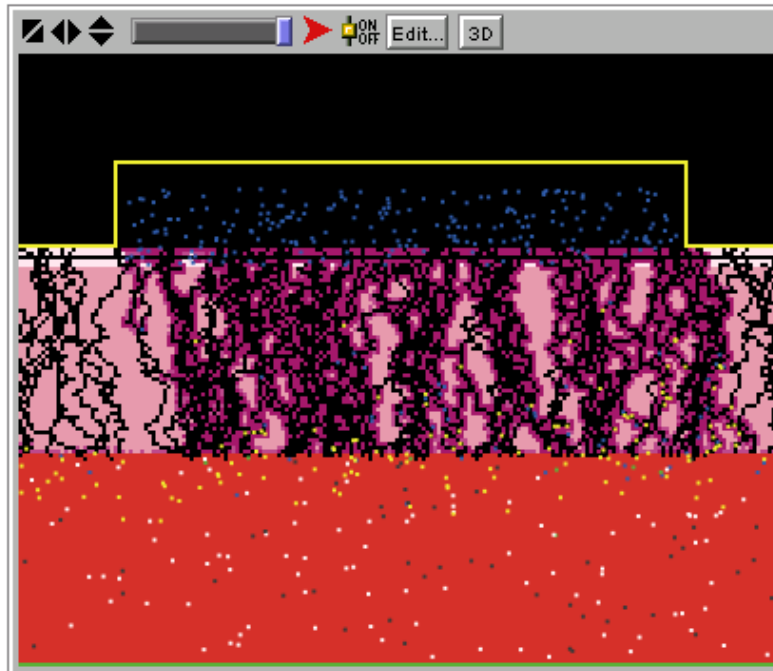
number-of-nonreactants 50

temp-scale 5




count ionized 69 dosage * .7

count nicotine 249 dosage * .3

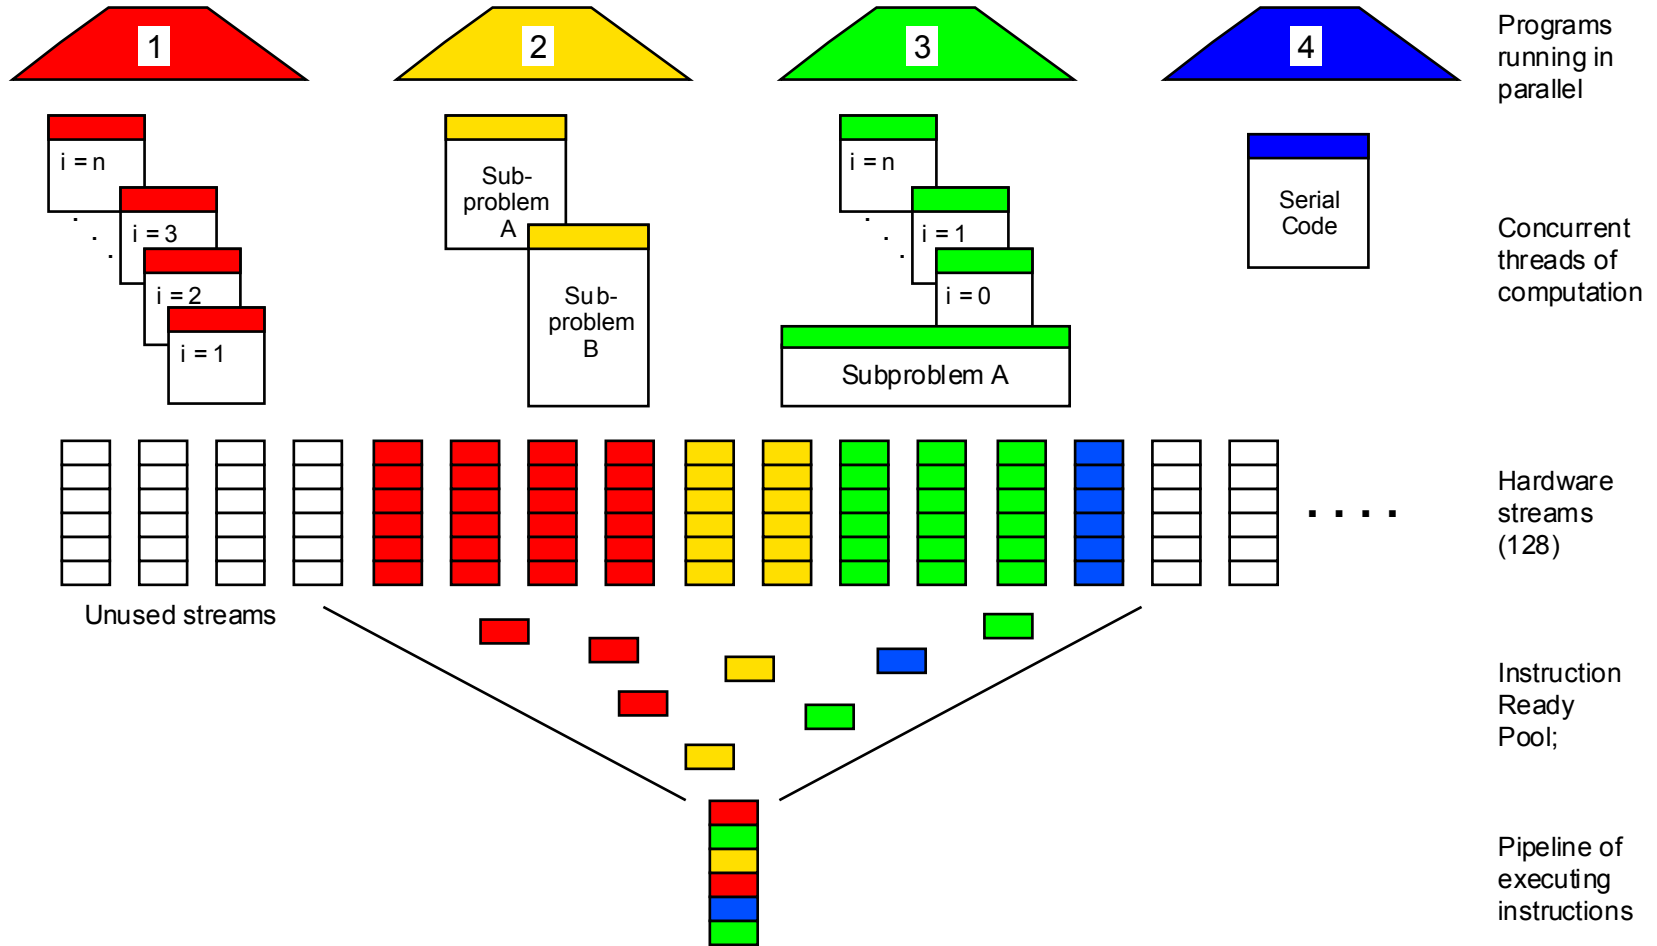
count stuck 16 dosage * .08



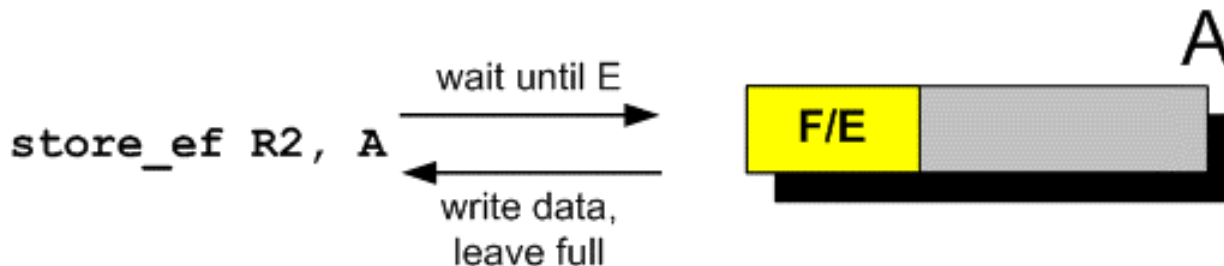

- Molecule Color Key
- Blue = Nicotine
 - Yellow = Ionized Nicotine
 - Black = Obstacle
 - White = Cell Receptor Site
 - Green = Received Receptor Site
- Patch Color Key
- Yellow = Edge of Patch
 - Pink = Skin Tissue
 - Black = Open Pathway
 - Light Purple = Irritated Tissue
 - Dark Purple = Irritated Tissue 2
 - Orange = Nicotine Absorption Site
 - Red = Bloodstream



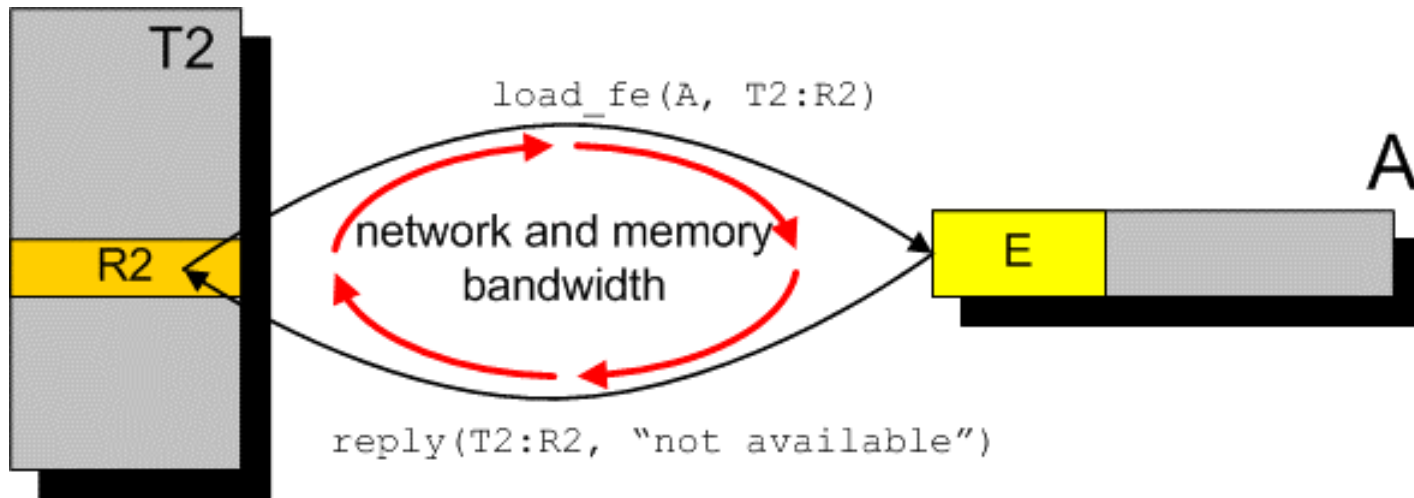
- ◆ Very lightweight; managed by compiler and runtime
 - Thread creation, termination, synchronization, and scheduling are performed by user space library code
- ◆ Runtime requests OS to add processors
- ◆ Thread virtualization
 - Application is coded without knowledge of stream count
 - Runtime migrates thread state between memory and stream
- ◆ Implicit creation by compiler and user directives
- ◆ Explicit creation using future statement




- ◆ F/E status of a location influences load/store behavior and loads/stores can modify F/E status.
 - Loads/stores can synchronize on F/E.
 - Cheap and abundant synchronization operations.
 - Prior MTA experience implementing f/e semantics.



- ◆ What happens if A is not in the required state when the `load_fe` or `store_ef` arrives?
 - Cray MTA f/e behavior for `load_fe` with A empty.



- Eventually a timeout occurs, the thread traps to a software handler; the thread state is saved to memory and the thread sleeps until A is marked full.
- When A becomes full, thread state is reloaded from memory and the `load_fe` is restarted.



```

typedef struct {
    int value;
    int n_edges_in;
    union {
        int *edges_in;
        int streamid;
    } u;
    int n_edges_out;
    int *edges_out;
} Node;
  
```

```

typedef struct {
    int from;
    int to;
    int weight;
    int value;
} Edge;
  
```

- ◆ Dataflow execution of input correlated streams
 - weighted sum at each node
- ◆ Nodes, edges, streams are arrays
 - MTA randomly scatters data in physical memory

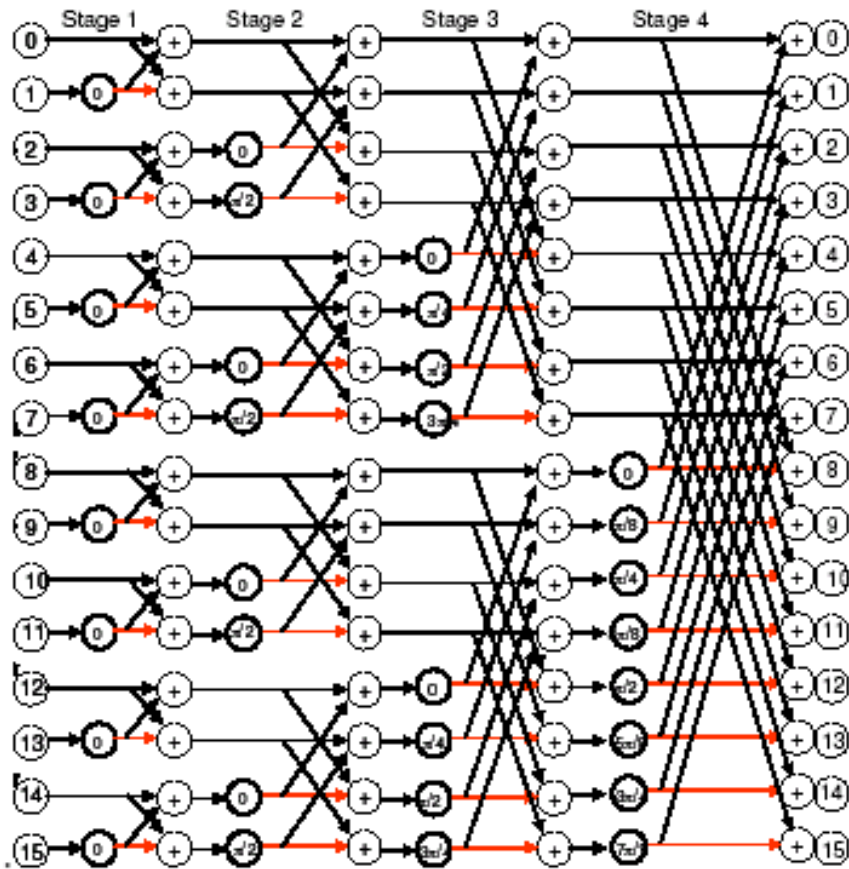
```
for each data set
  for each stage of the graph
    #pragma mta assert parallel
    for each node in that stage
      read()
      compute()
      write()
```

```
for each data set
  #pragma mta assert parallel
  for each node in graph
    readfe()
    compute()
    writeef()
```

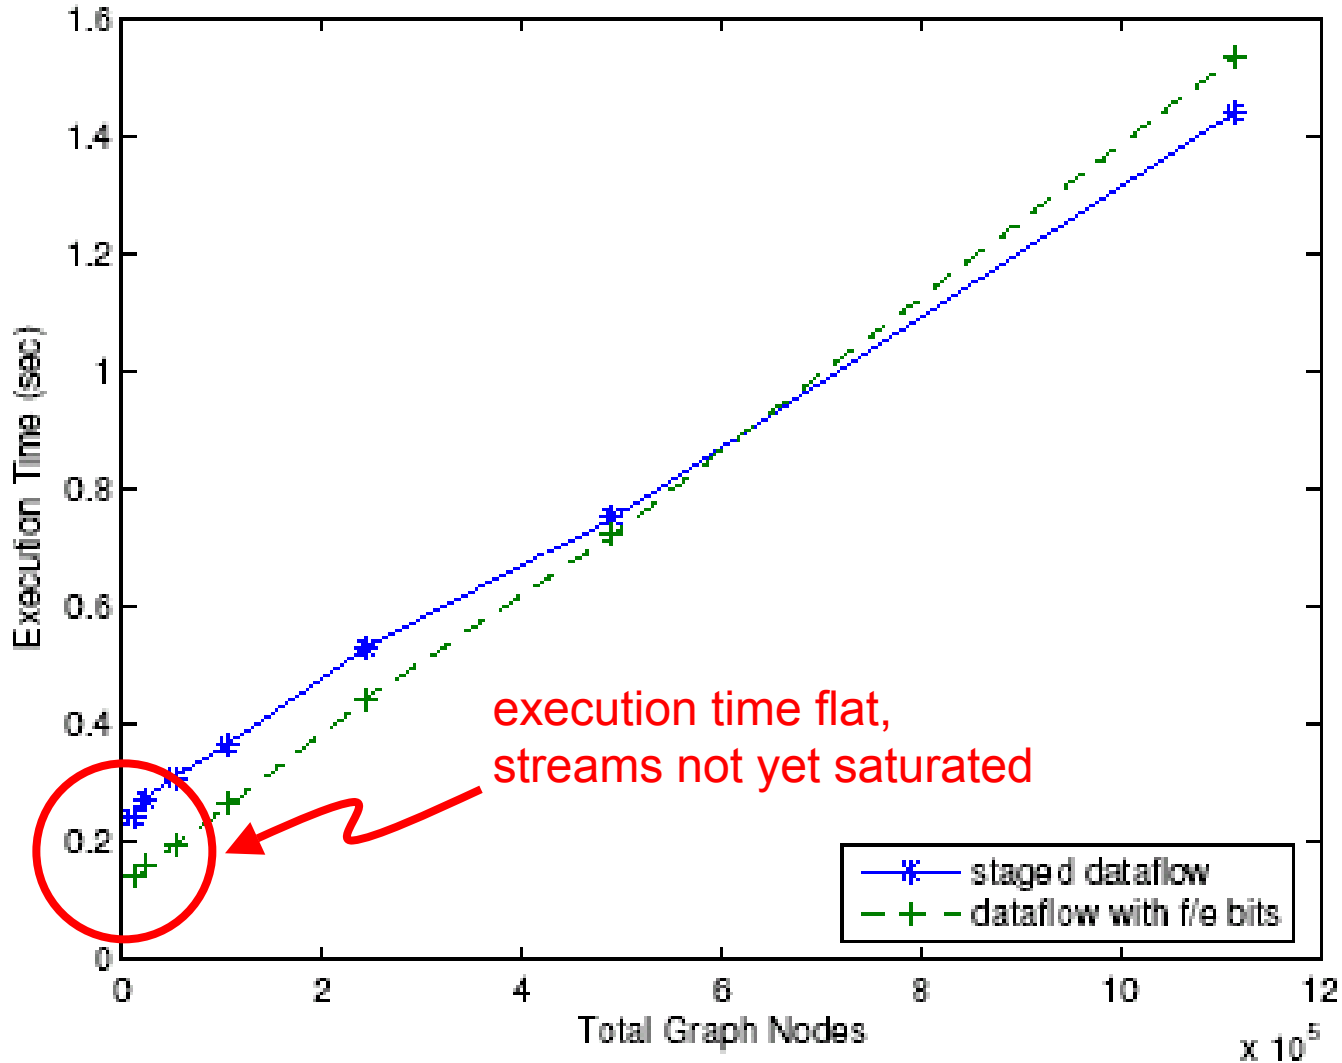
- ◆ Synchronous stages
 - stages can be hard to identify in random graphs
- ◆ Asynchronous using full/empty bits
 - no need to identify stages





Experiment: Parameterized Butterfly Graphs

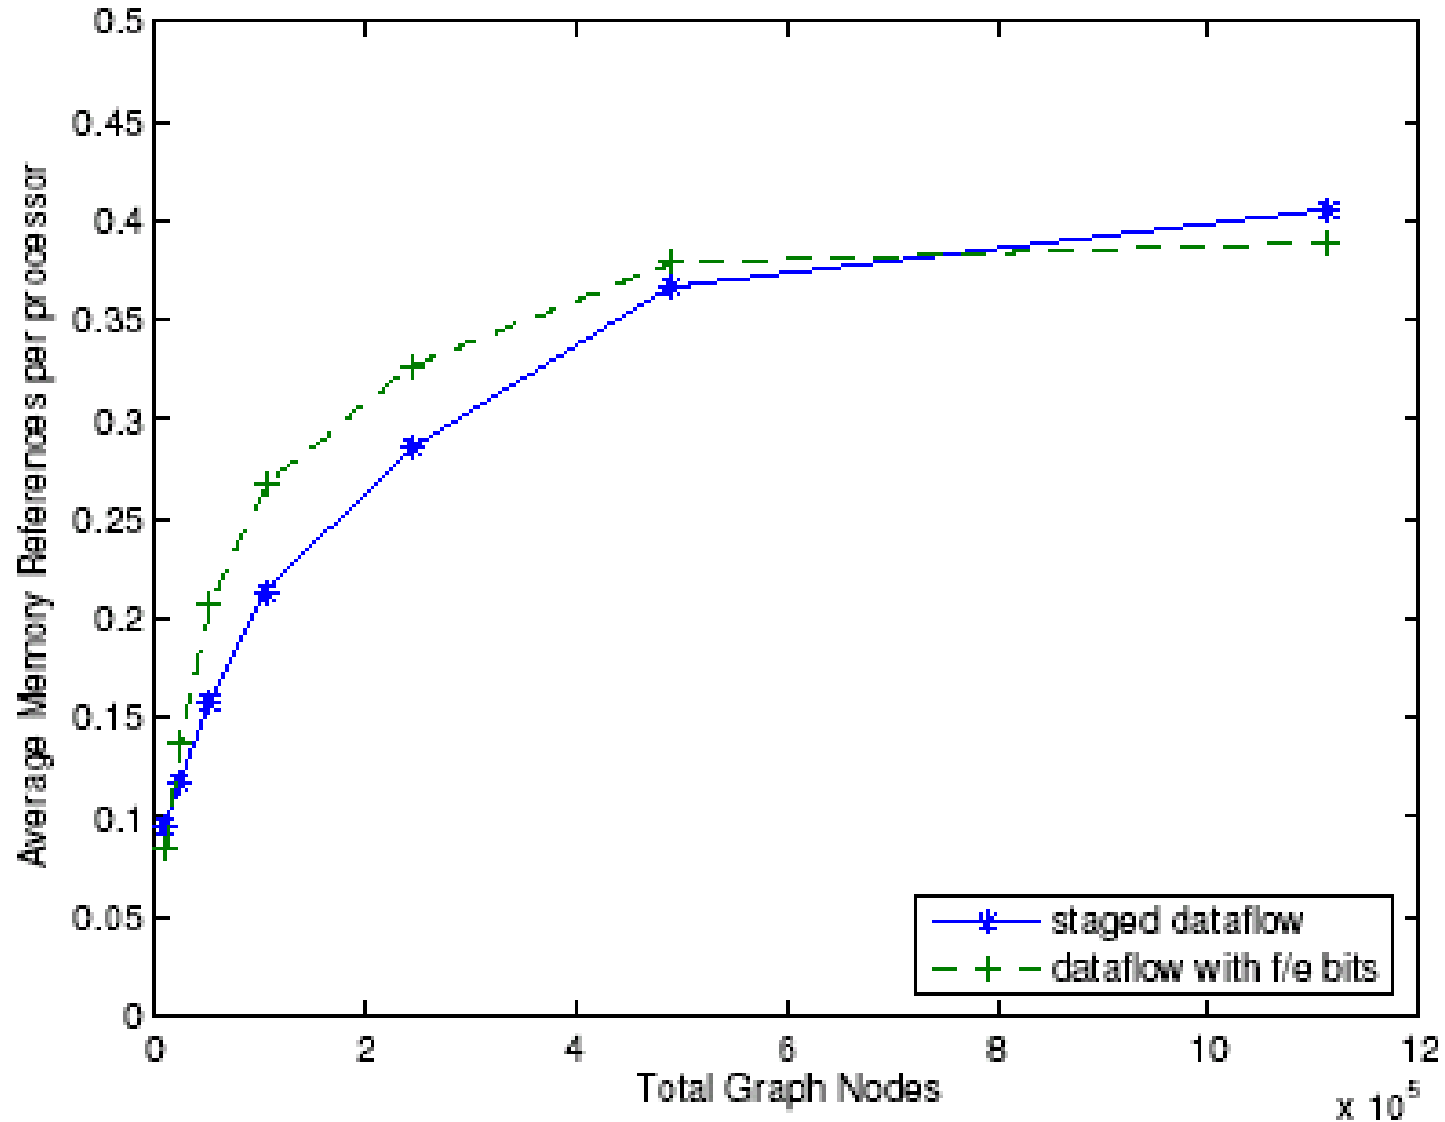


- ◆ Varied
 - Number of nodes
 - Radix
 - Randomly deleted edges
- ◆ Run on 10 MTA-2 processor nodes









```
#pragma mta assert parallel
for each node in graph
  for each dataset
    readfe()
    compute()
    writeef()
```

- ◆ Deadlock after 512 nodes!!!
- ◆ Why?
 - Outstripped number of available streams
 - Runtime scheduling becomes an issue

- ◆ Active graphs are a natural way to formalize systems of interacting agents
 - Easy implementation on MTA
- ◆ Very low overhead synchronization using full/empty bits
 - Avoids need for explicit scheduling—to a point
- ◆ What's next?
 - More applications, including ZCHAFF SAT solver
 - Investigate PGAS programming models and performance on Eldorado—path to Cascade