

Parallel Computation

"Faster" is the only reason

But...

greater programming complexity

new kinds of bugs

...and not much help for fixing them.

Parallel Computation

"Faster" is the only reason

But...

greater programming complexity

new kinds of bugs

...and not much help for fixing them.

Can the day or week of user effort be recovered?

Parallel Computation

"Faster" is the only reason

But...

greater programming complexity

new kinds of bugs

...and not much help for fixing them.

Can the day or week of user effort be recovered?

16384 core EPFL IBM BlueGene/P

1 hour at 850MHz

6 months at 3GHz

Parallel Computation

A simulation run takes about a second
want to do 1000's of them,
varying a dozen or so parameters.

A simulation run takes hours.
want to spread the problem over several machines.

Parallel Computation

A simulation run takes about a second
want to do 1000's of them,
varying a dozen or so parameters.

- Screensaver Calin–Jageman and Katz, 2006
- Bulletin–board (Linda)

A simulation run takes hours.
want to spread the problem over several machines.

Parallel Computation

A simulation run takes hours.

want to spread the problem over several machines.

Network

Subnets on different machines

Cells communicate by:

logical spike events with significant
axonal, synaptic delay.

postsynaptic conductance depends
continuously on presynaptic voltage.

gap junctions

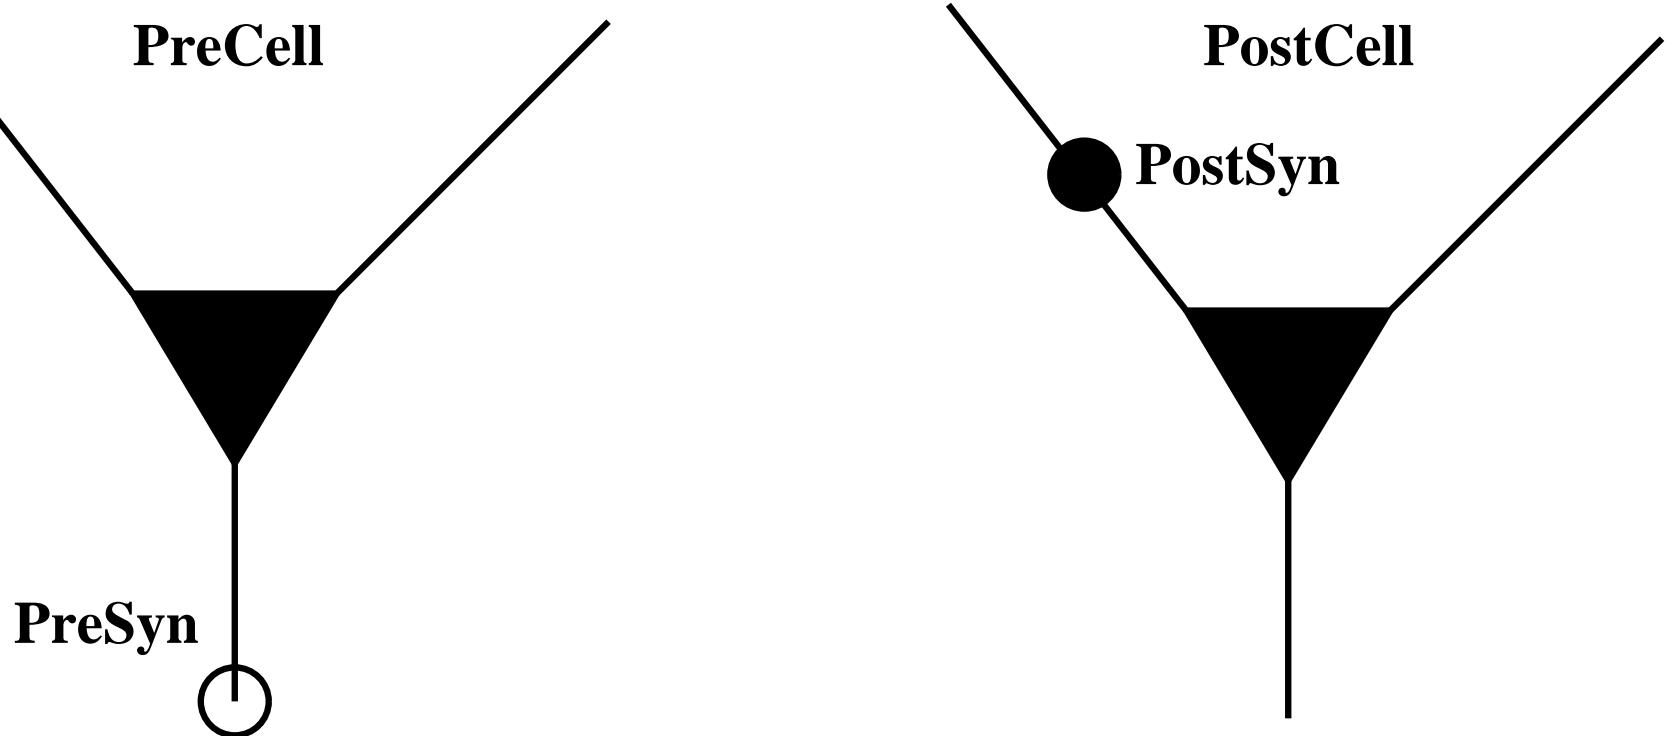
Parallel Computation

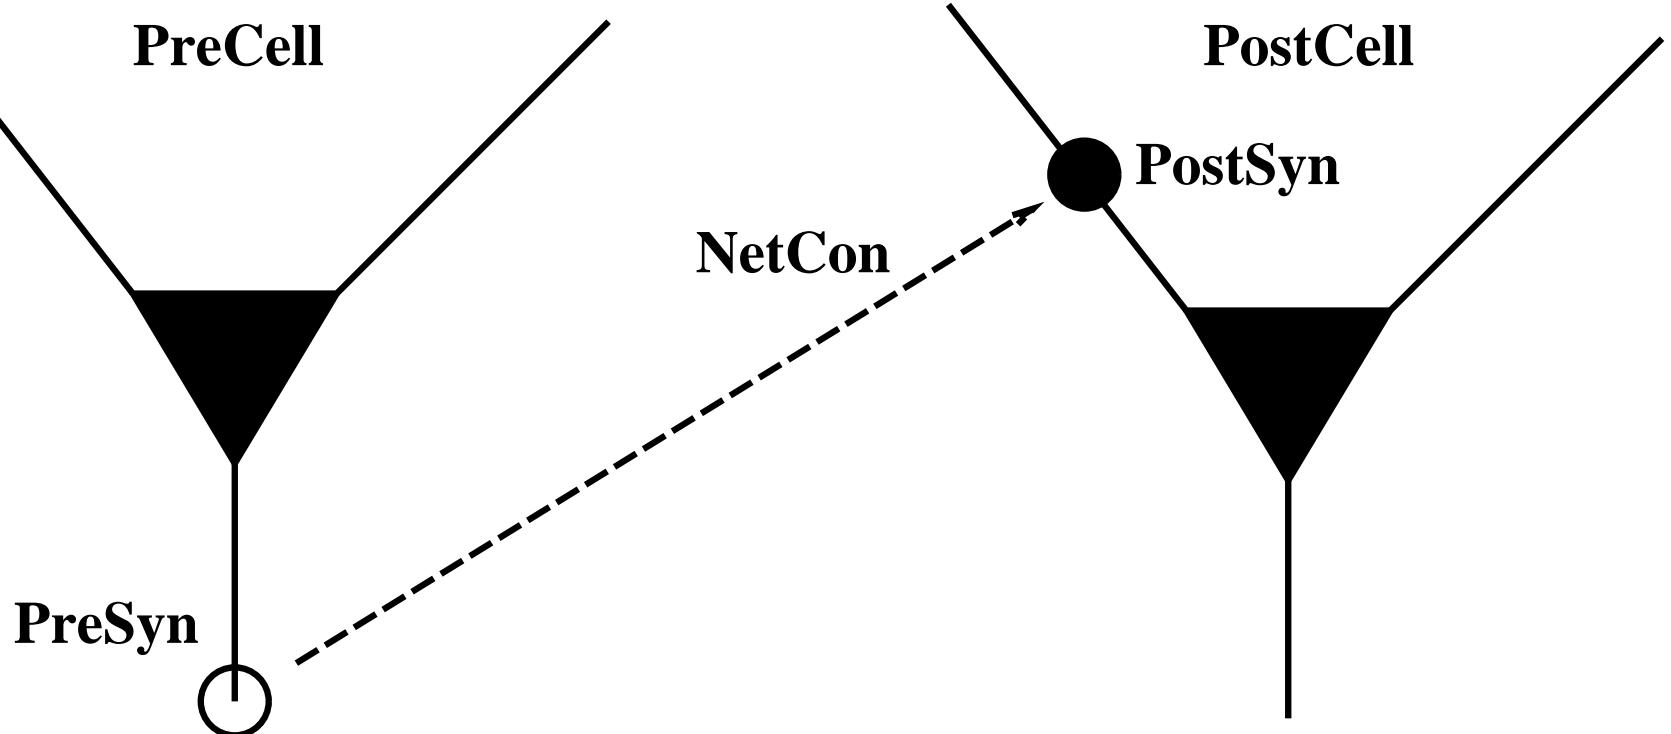
A simulation run takes hours.

want to spread the problem over several machines.

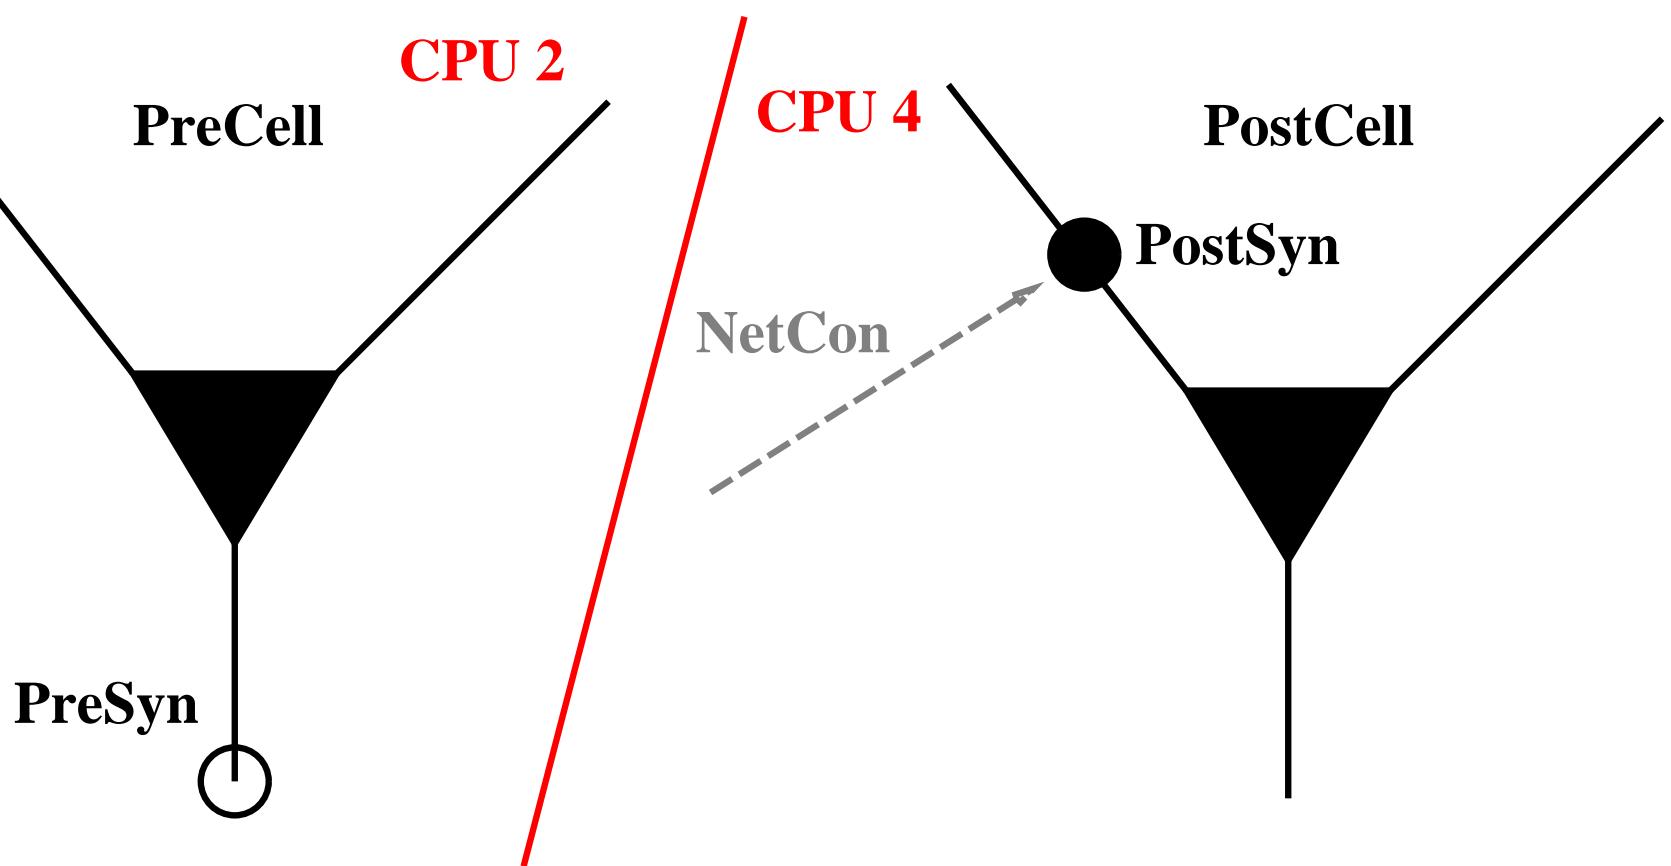
Single cells

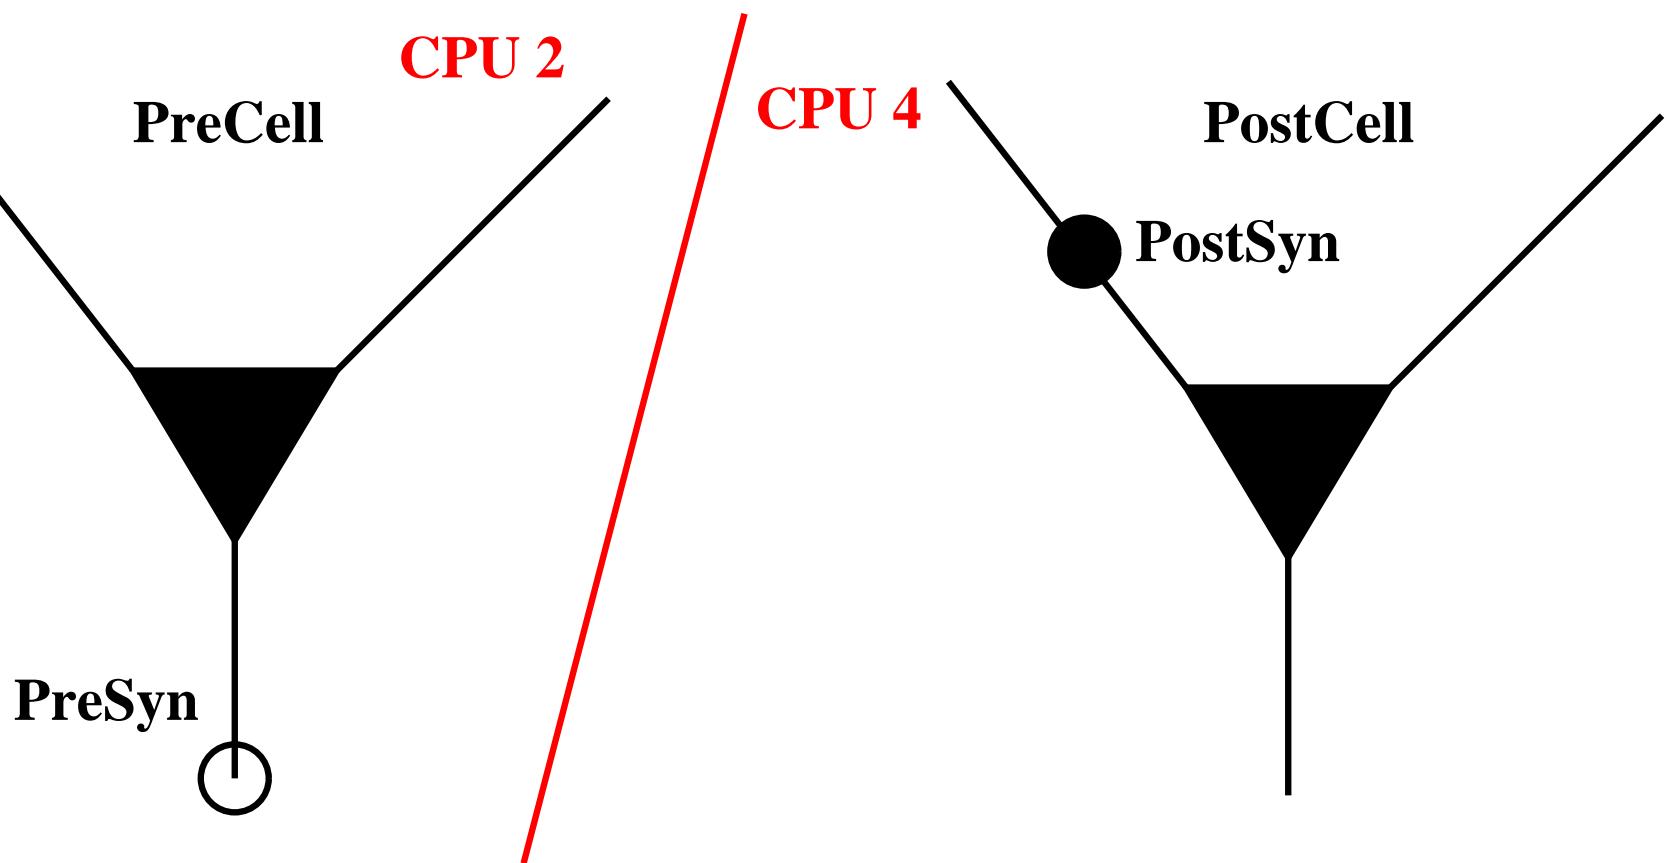
portions of the tree cable equation on
different machines.



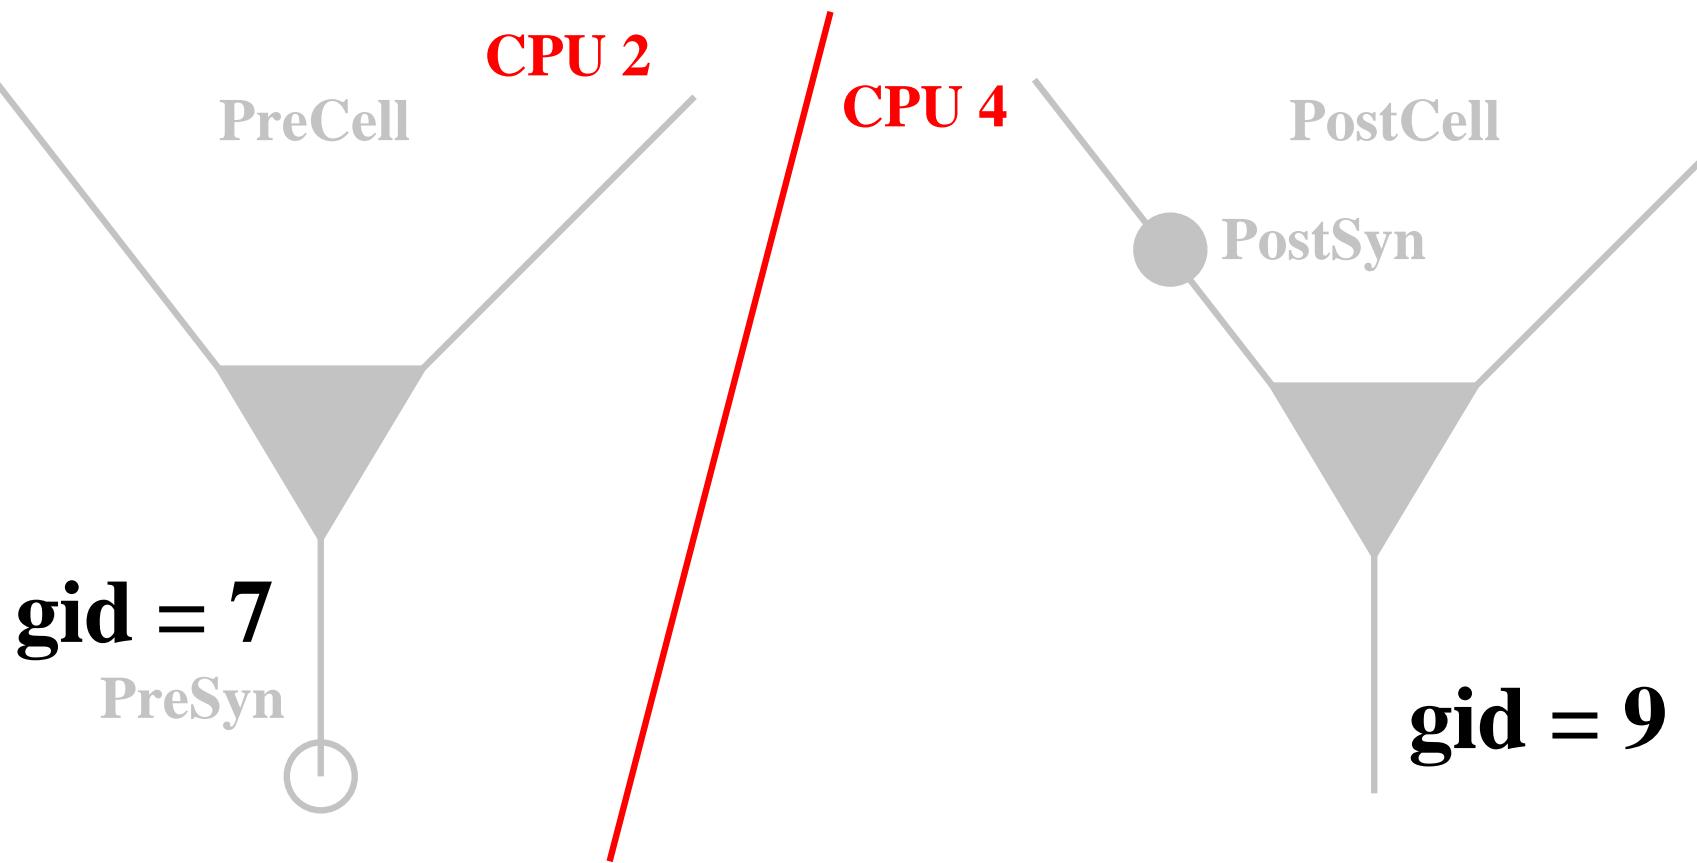


```
nc = new NetCon(PreSyn, PostSyn)
```





```
pc = new ParallelContext()
```



Every spike source (cell) must have a global id number.

CPU 0

pc.id 0
pc.nhost 5
ncell 14

•••

gid
0
5
10

CPU 3

pc.id 3
pc.nhost 5
ncell 14

gid
3
8
13

CPU 4

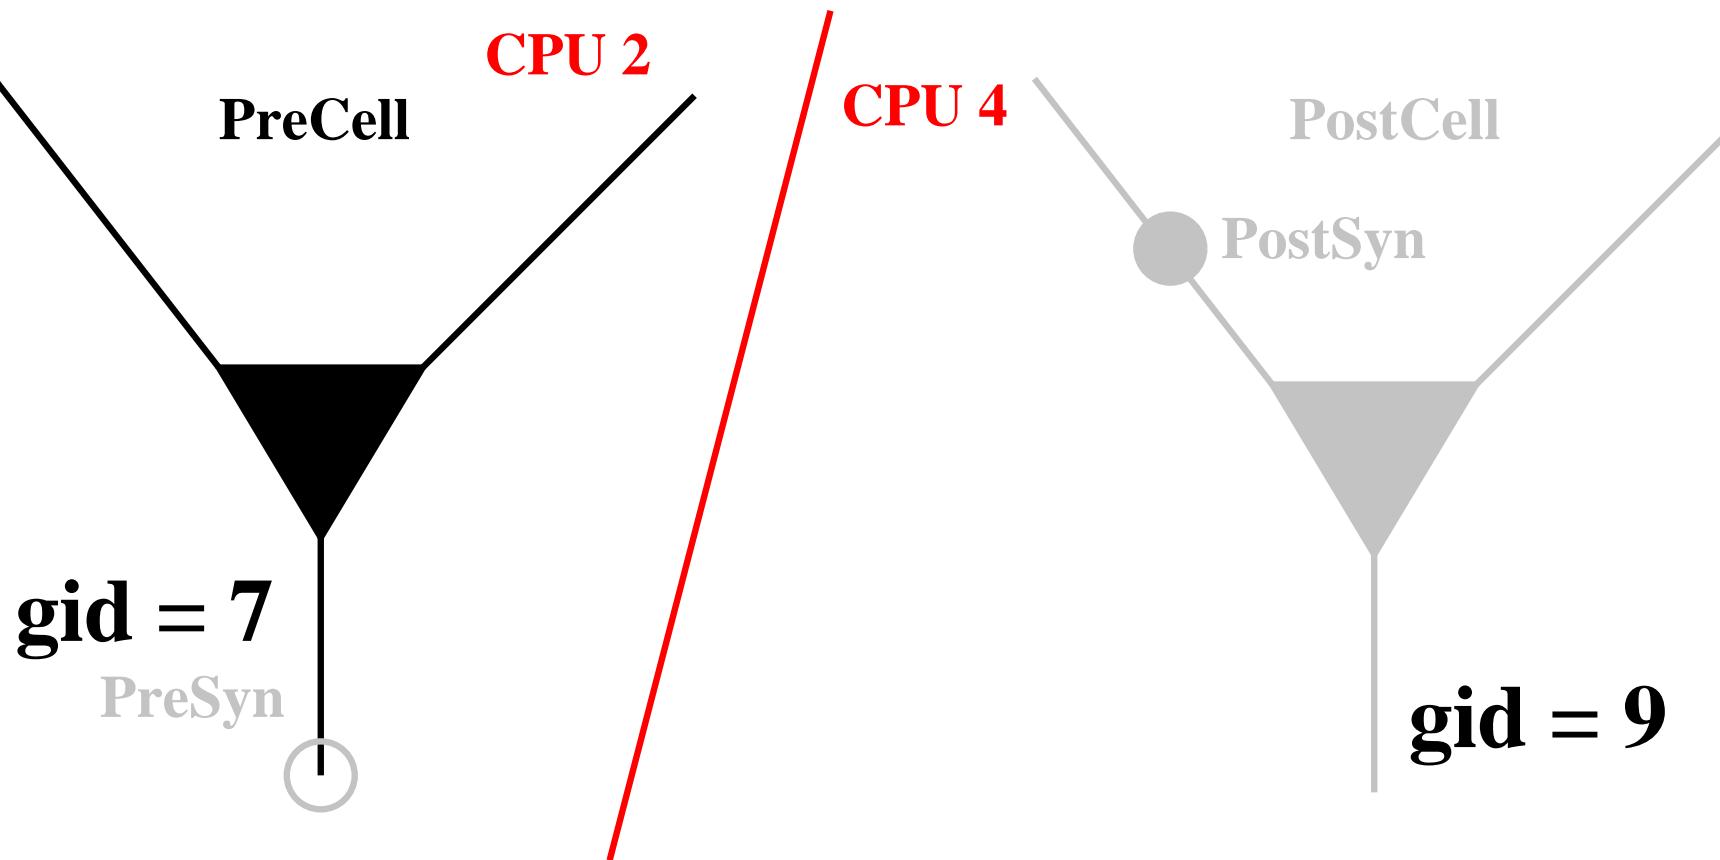
pc.id 4
pc.nhost 5
ncell 14

gid
4
9

An efficient way to distribute:

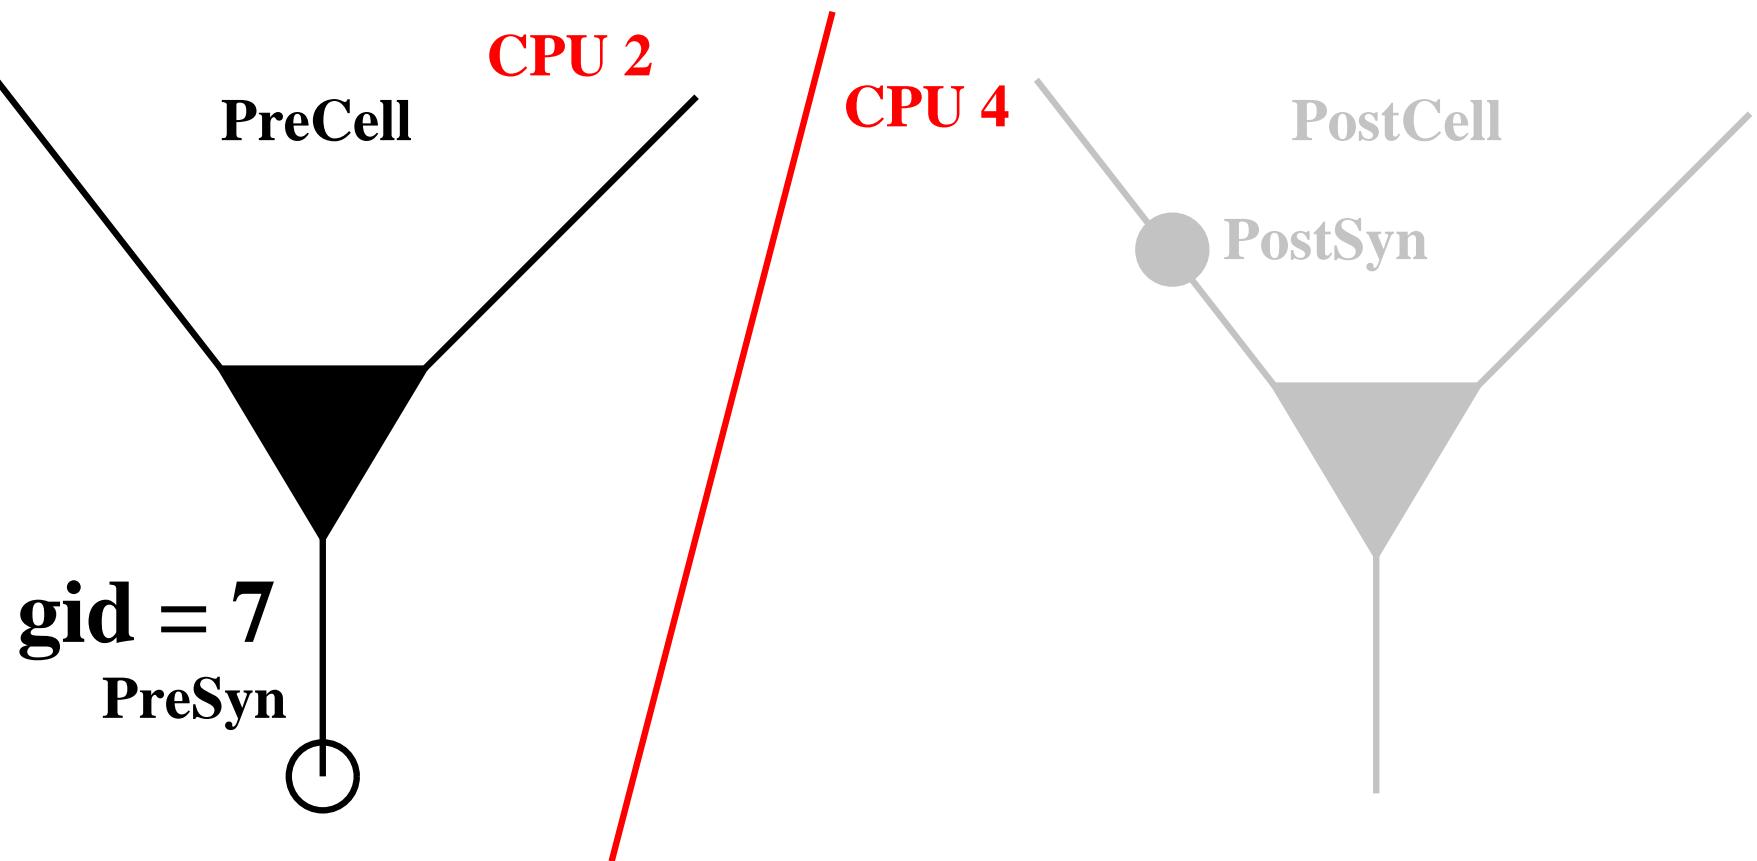
```
for (gid = pc.id; gid < ncell; gid += pc.nhost) {  
    pc.set_gid2node(gid, pc.id)  
    ...  
}
```

body executed only ncell/nhost times, not ncell.



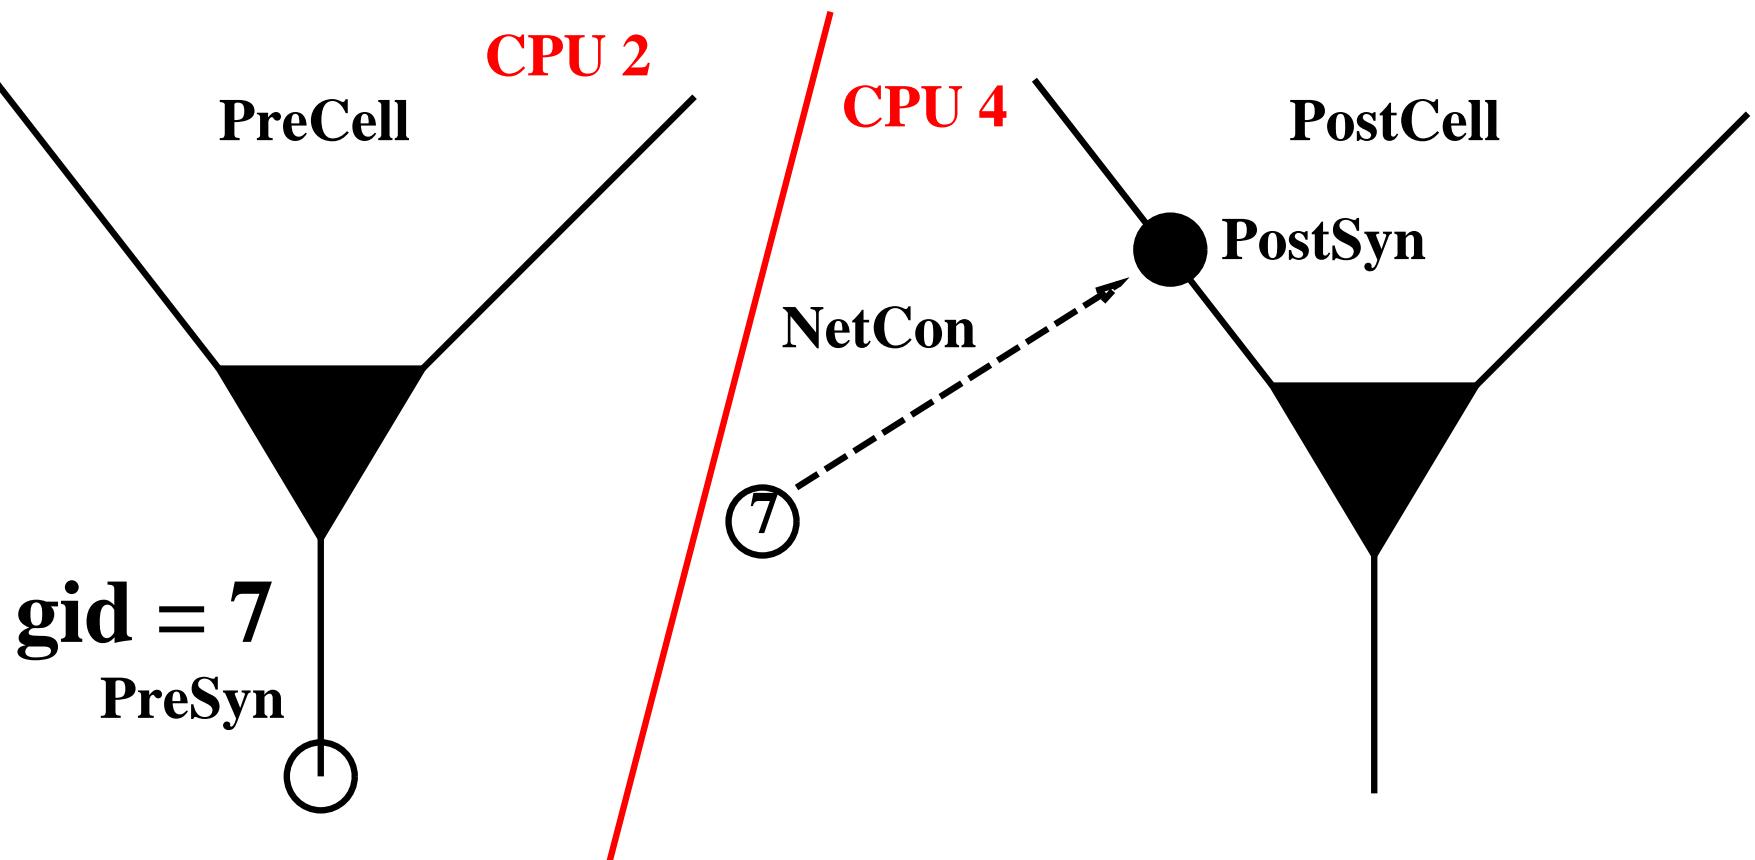
Create cell only where the gid exists.

```
if (pc.gid_exists(7)) {  
    PreCell = new Cell()  
}
```



Associate gid with spike source.

```
nc = new NetCon(PreSyn, nil)  
pc.cell(7, nc)
```



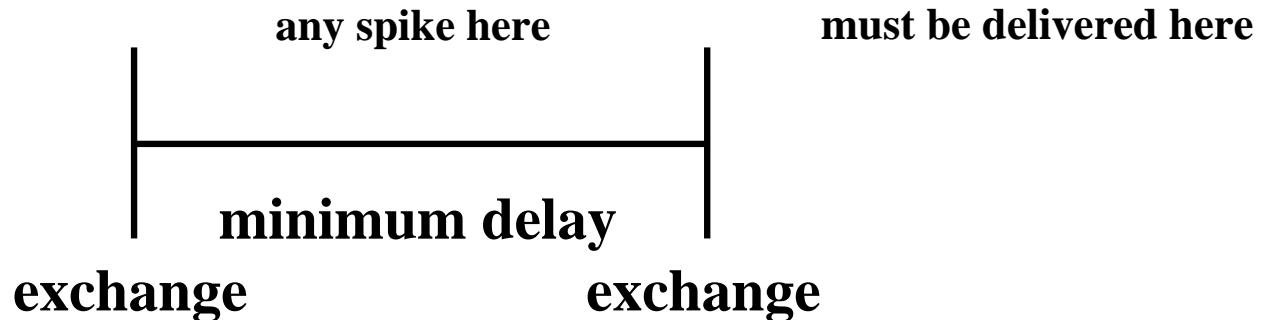
Create NetCon on CPU where target exists.

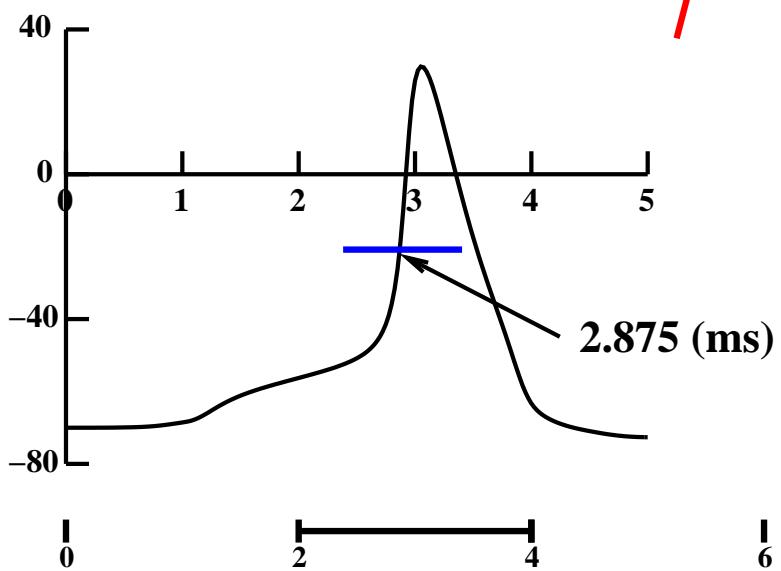
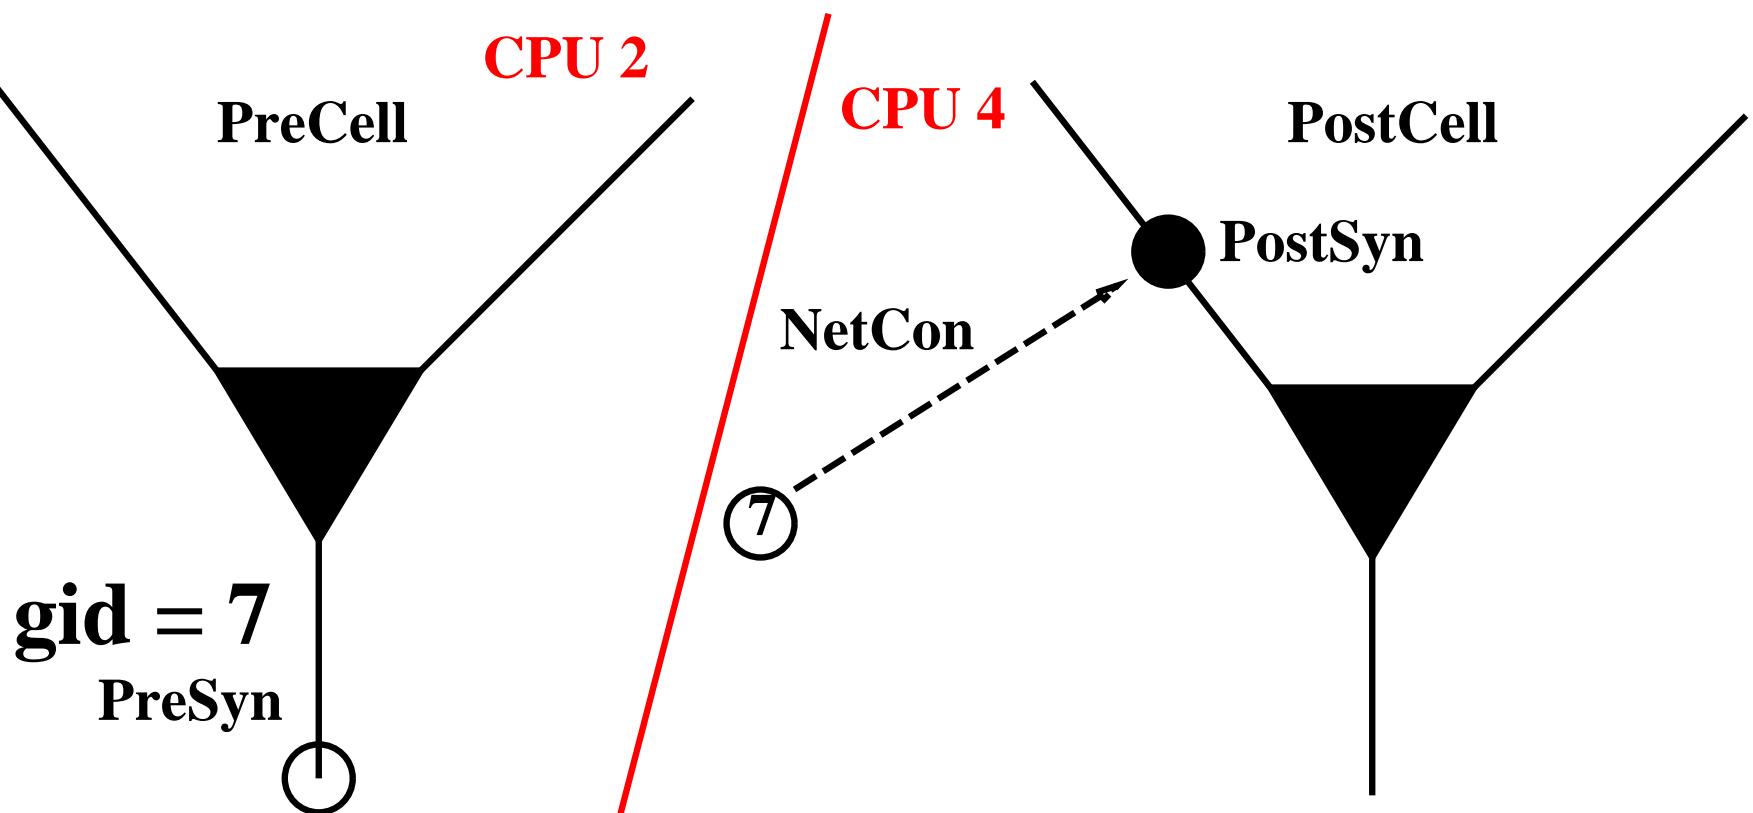
```
nc = pc.gid_connect( 7, PostSyn )
```

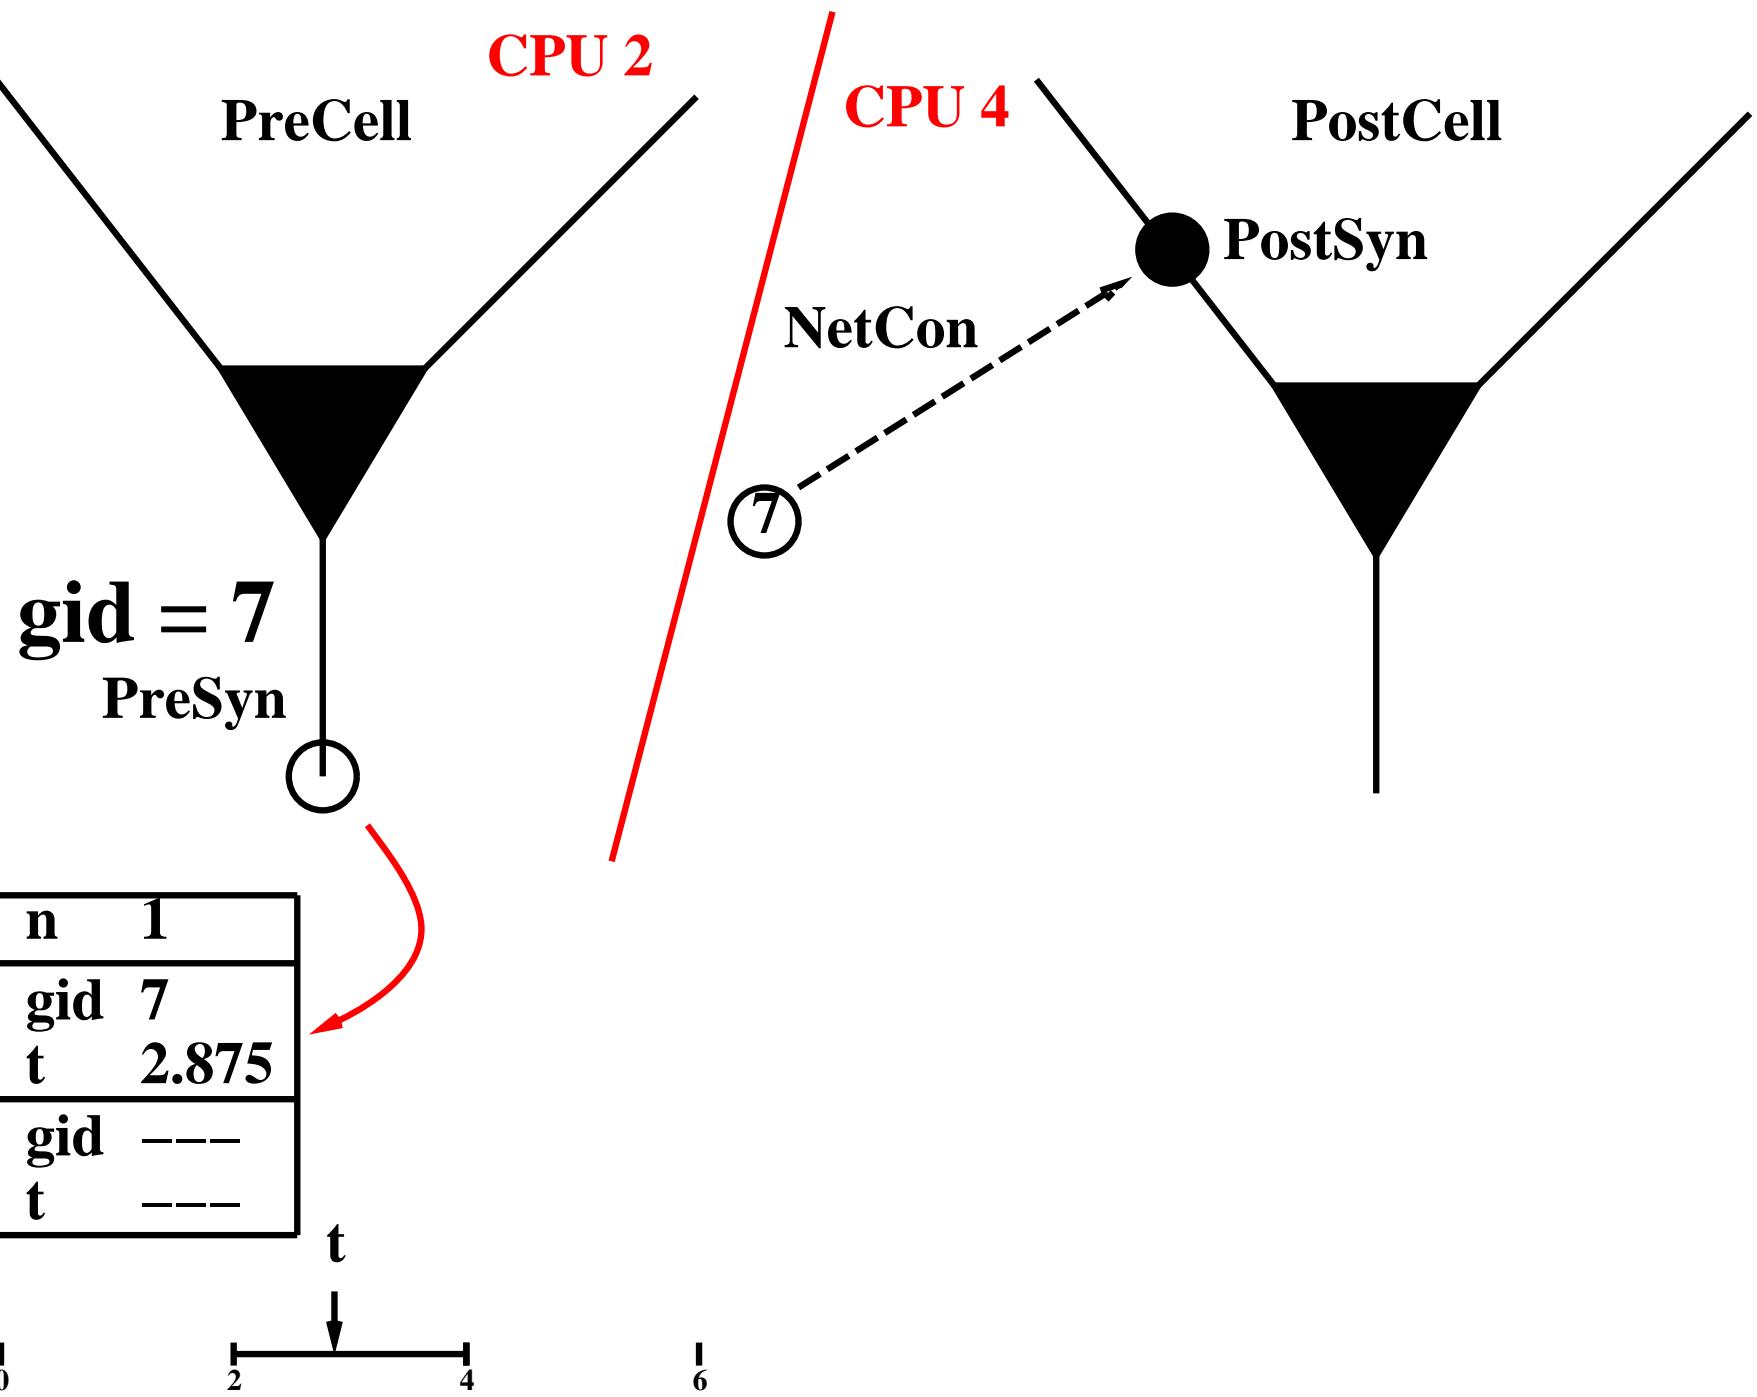
Run using the idiom

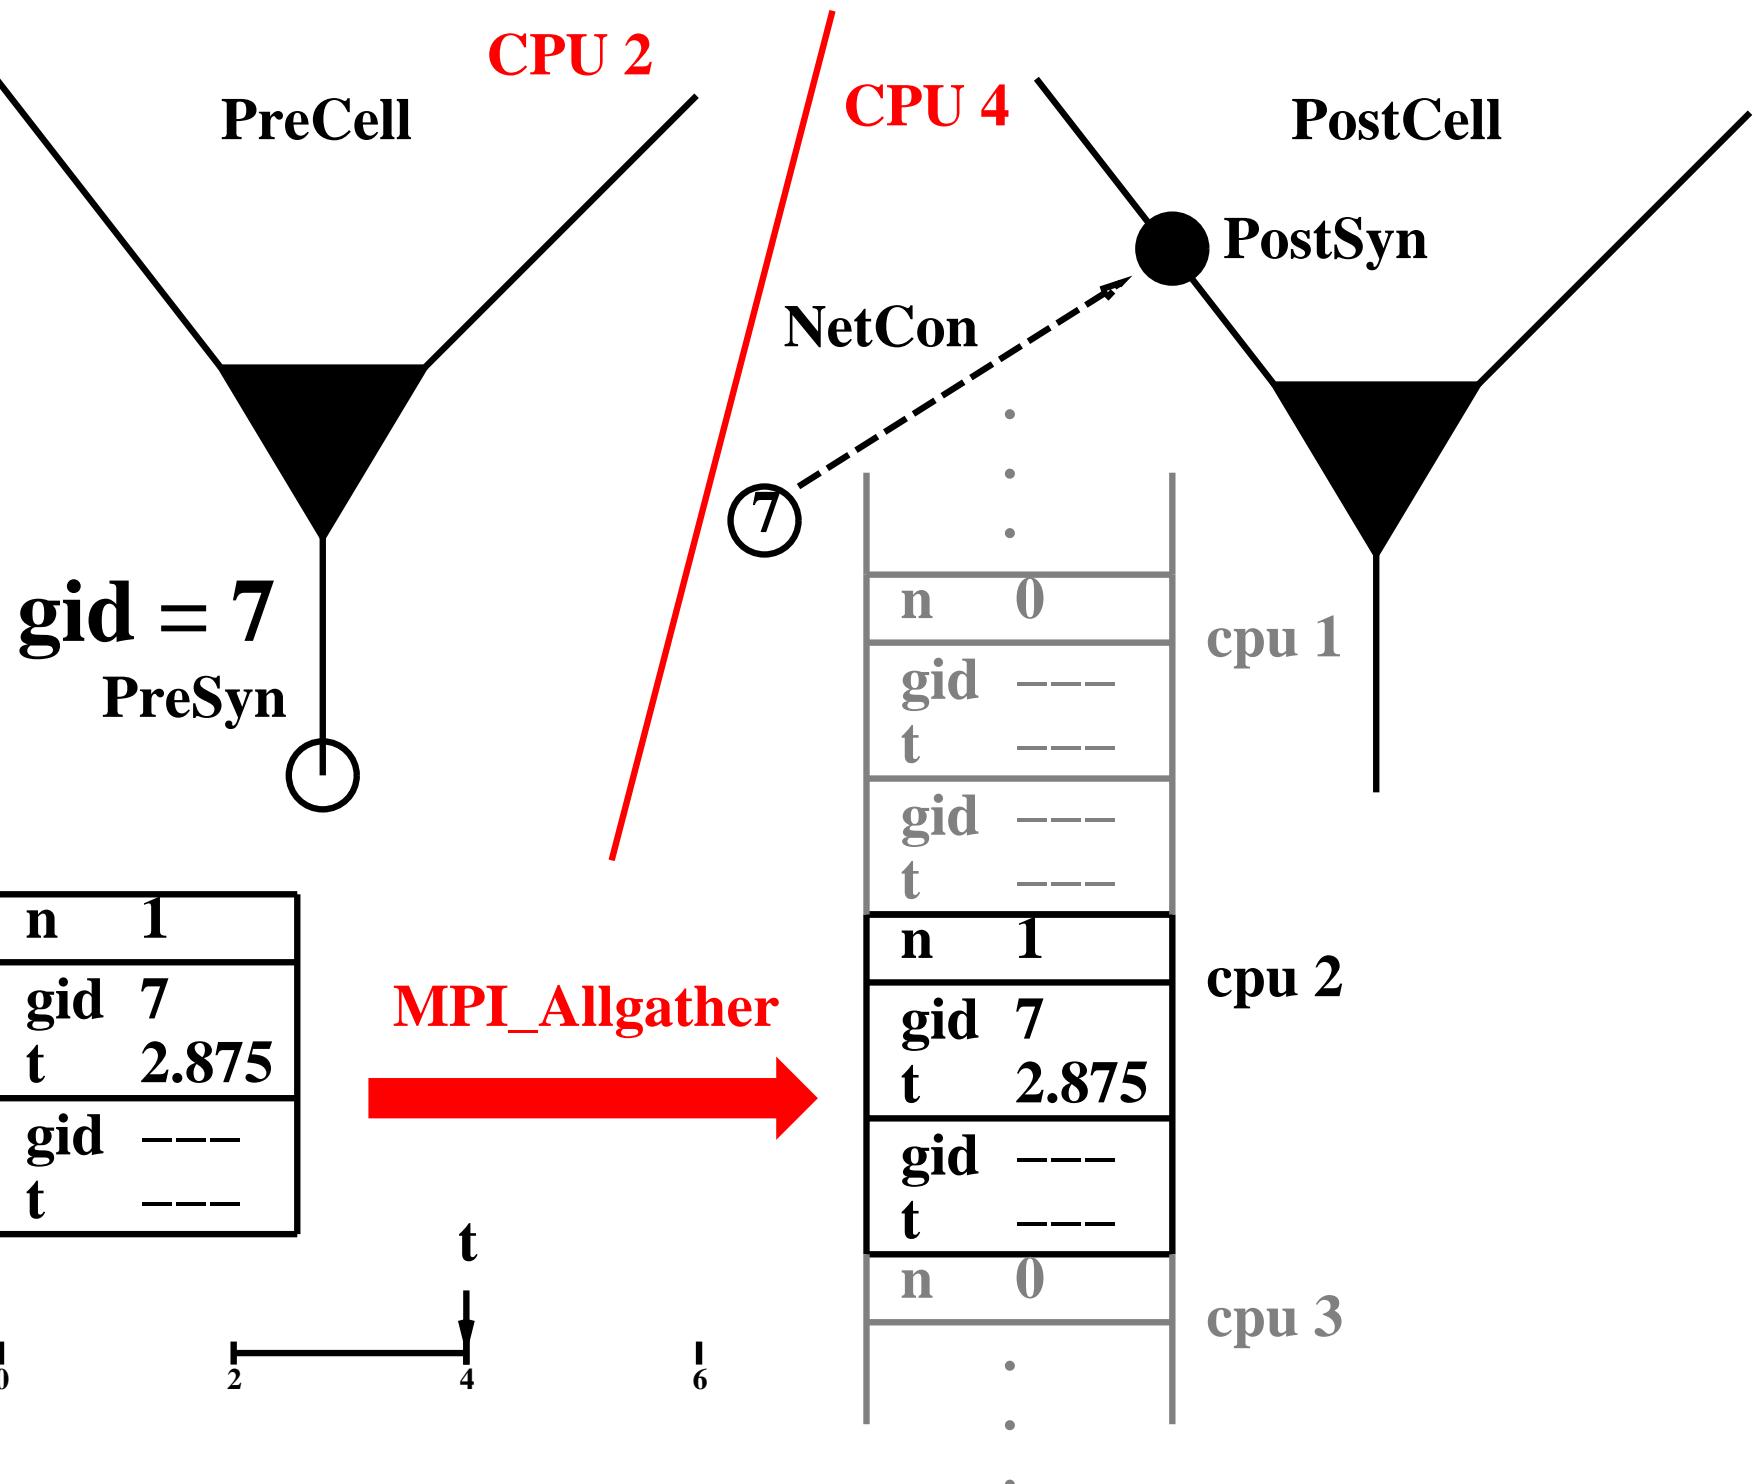
```
pc.set_maxstep(10)  
stdinit()  
pc.psolve(tstop)
```

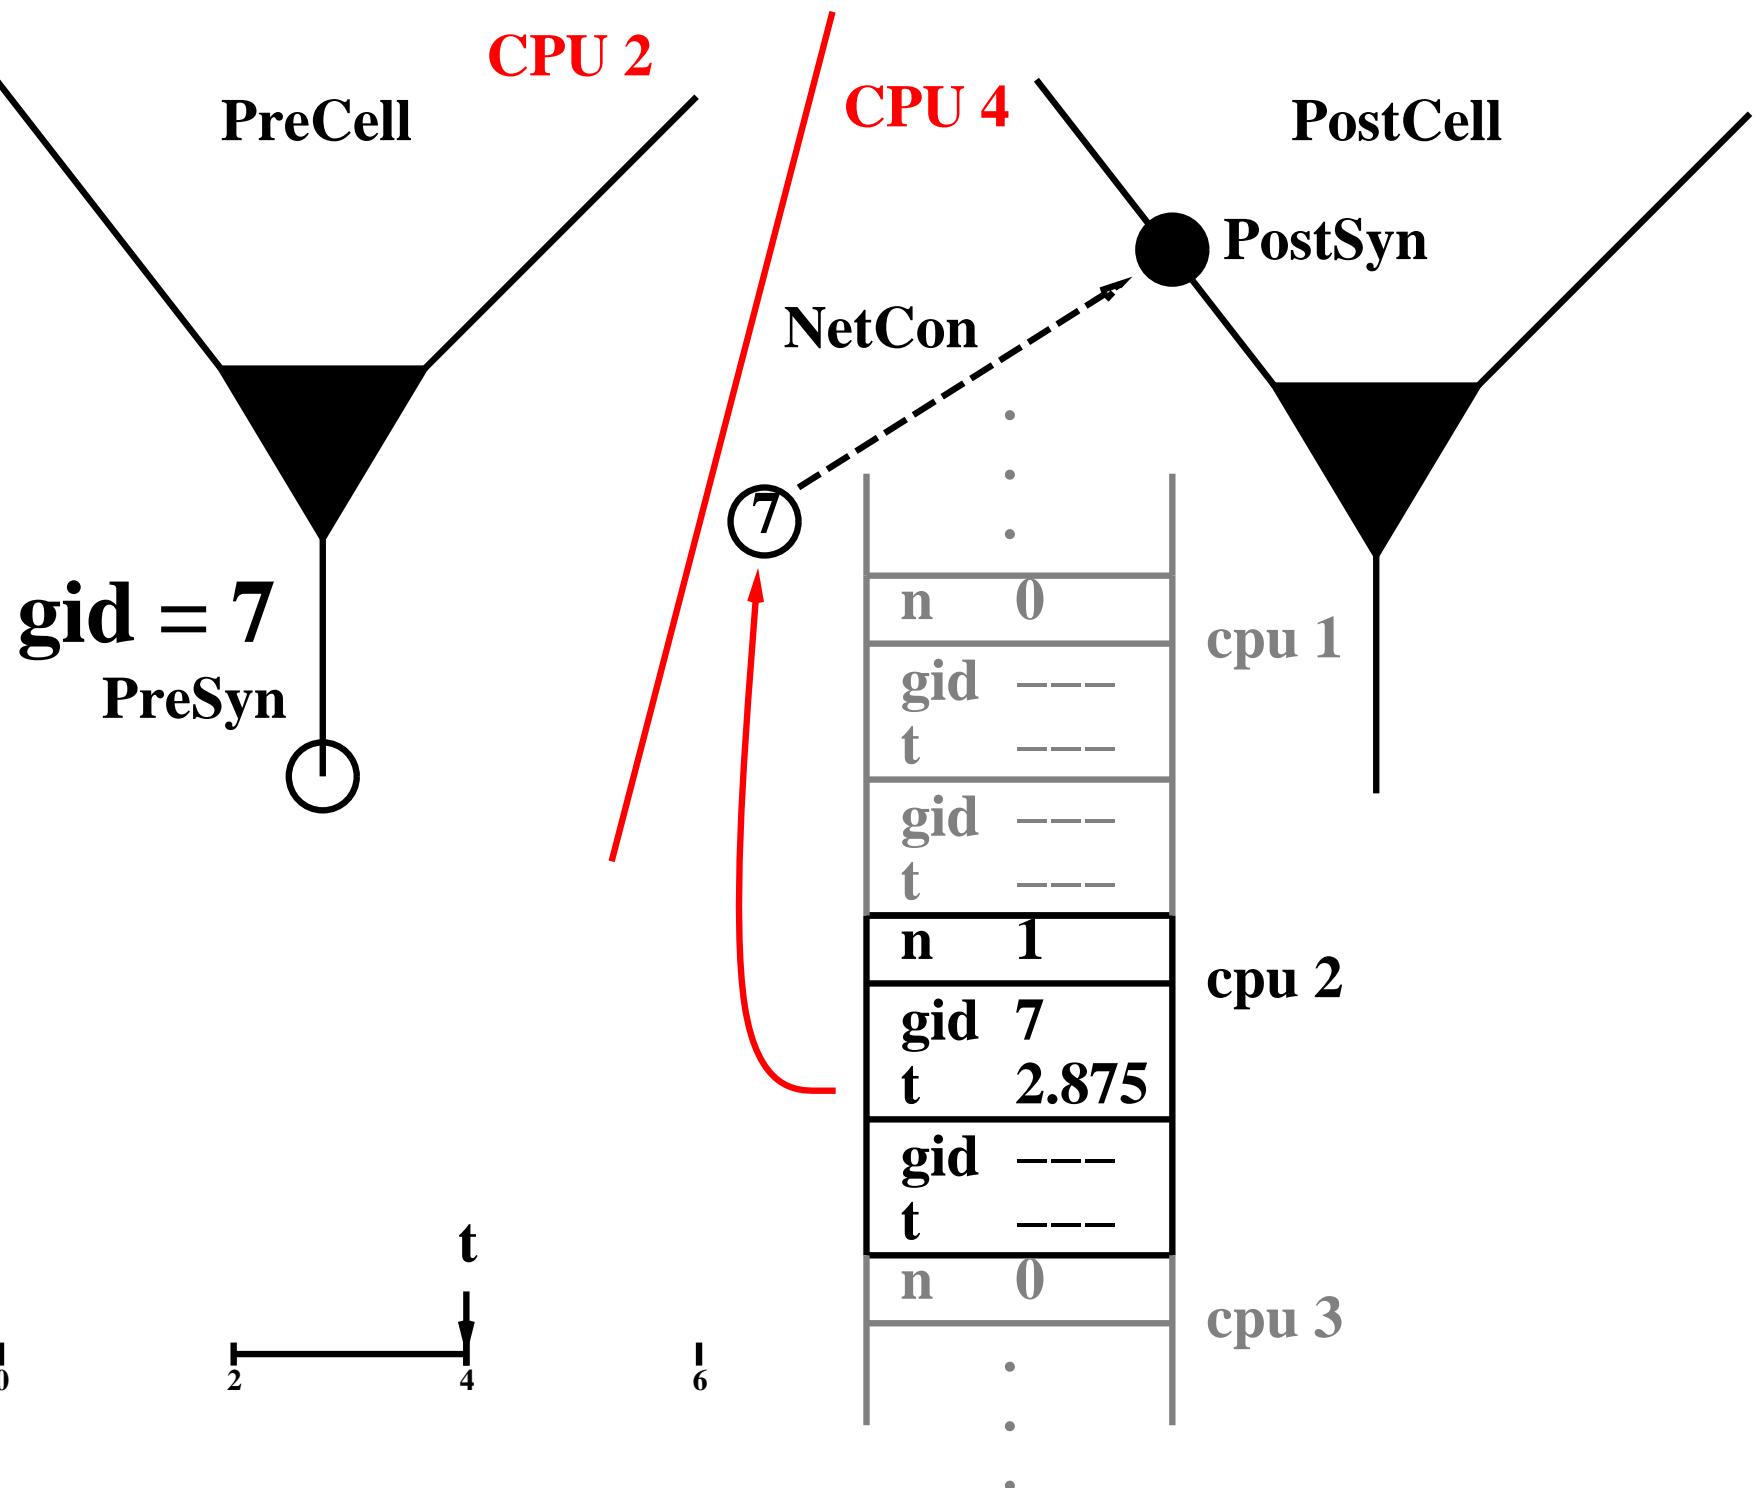
**pc.set_maxstep() uses
MPI_Allreduce
to determine minimum delay.**





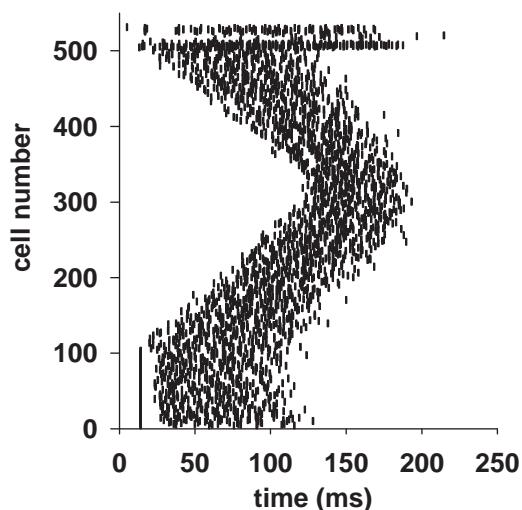




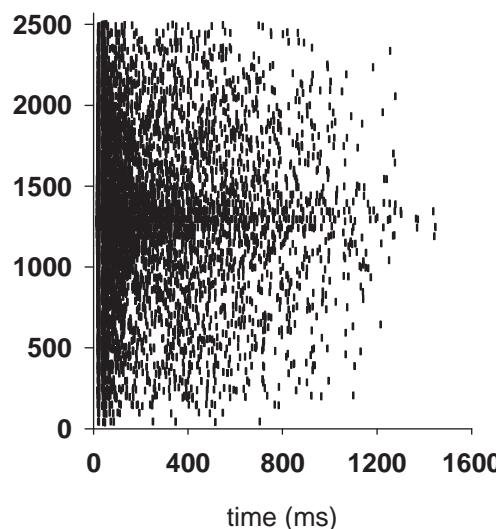


Migliore et al (2006) J. Comput. Neurosci. 21(2):119

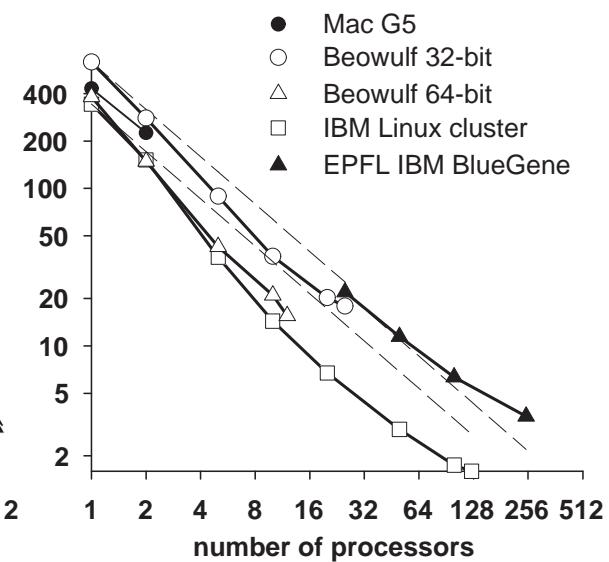
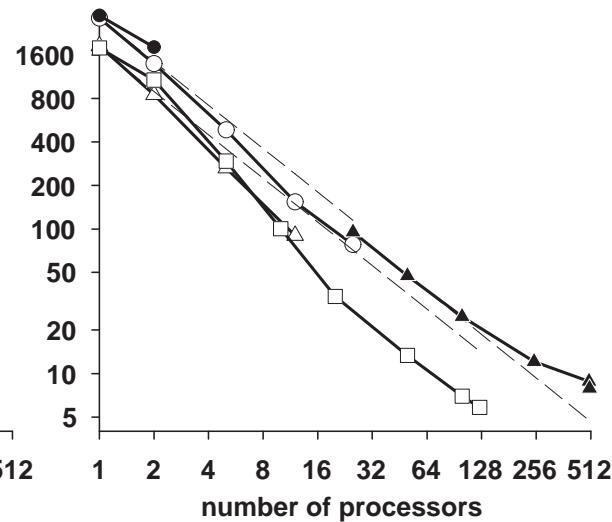
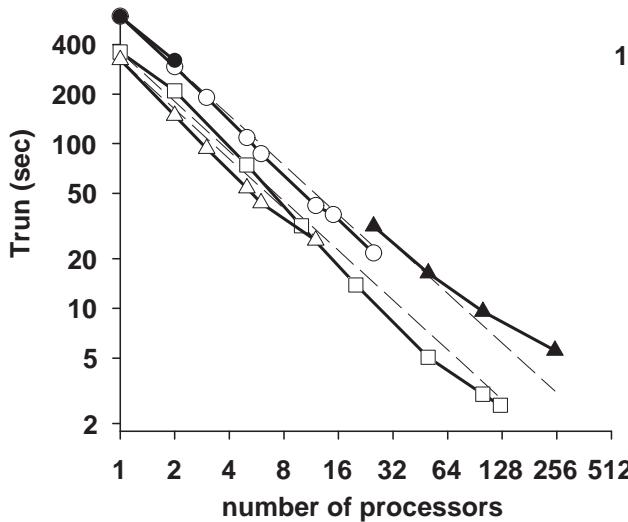
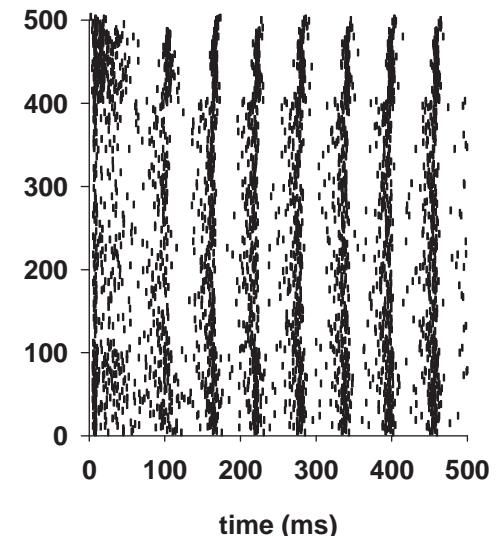
Santhakumar et al. (2005)



Davison et al., (2003)

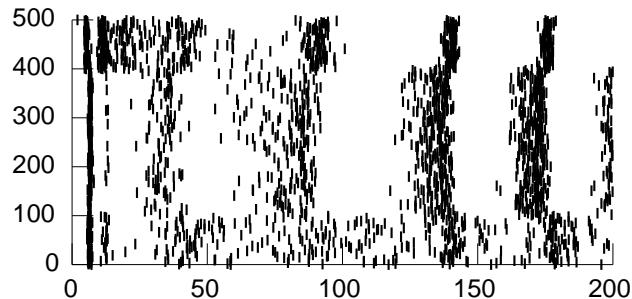
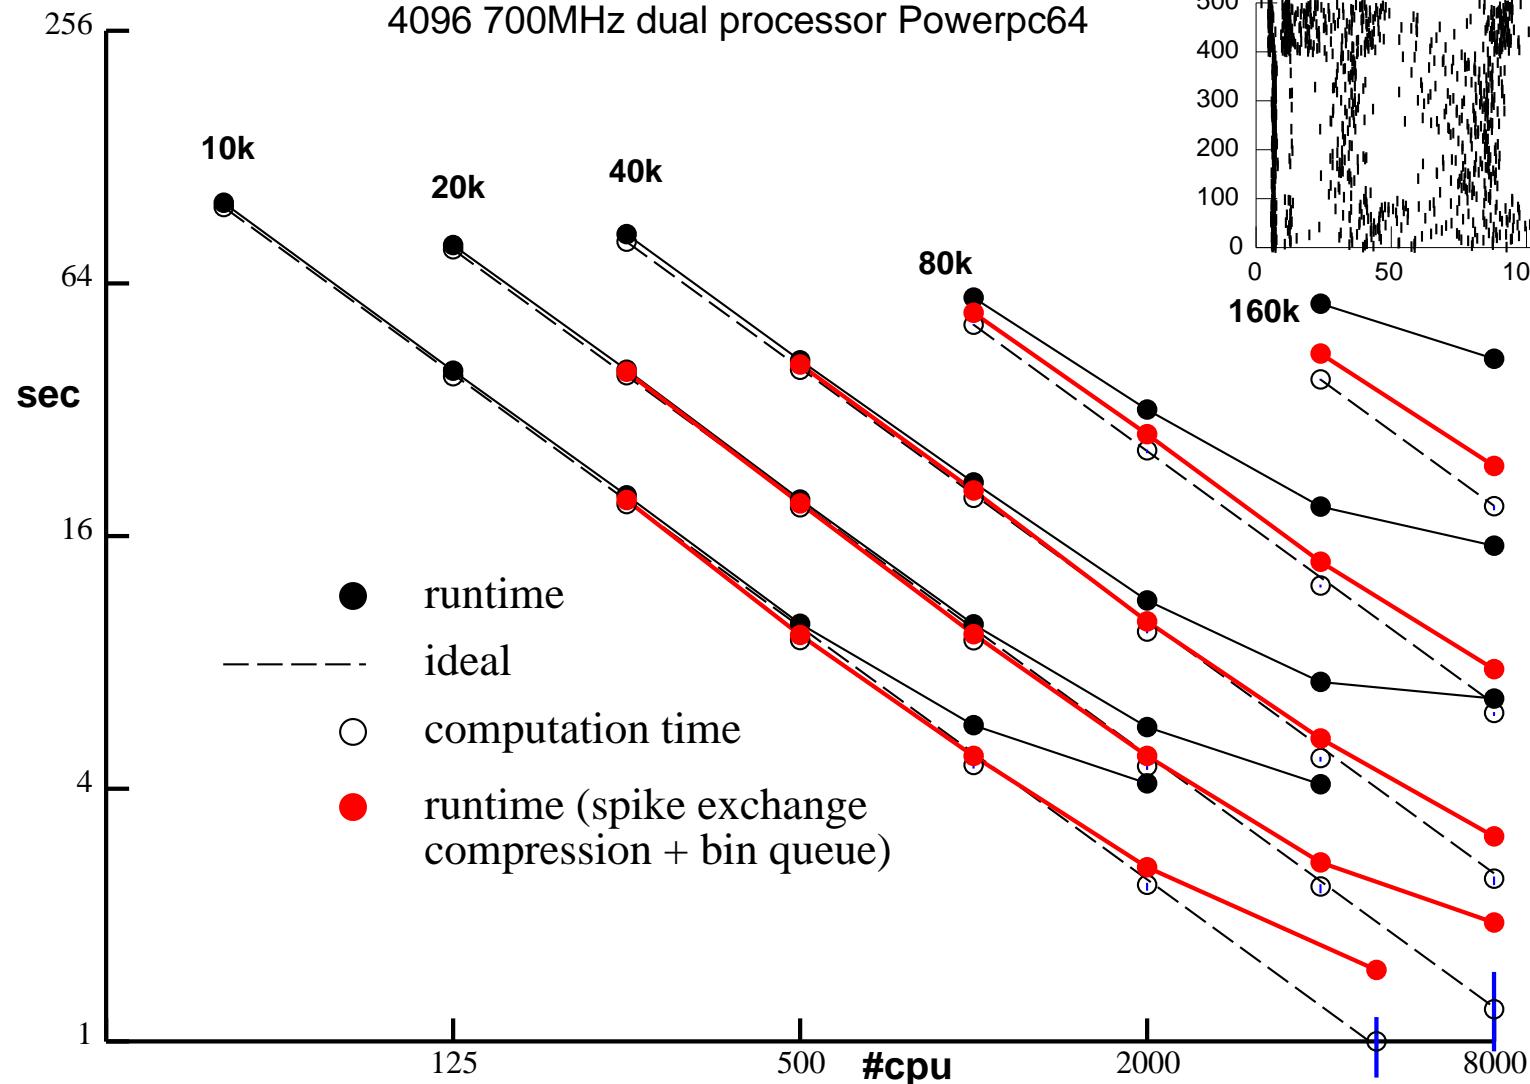


Bush et al., (1999)



EPFL IBM BlueGene/L

4096 700MHz dual processor Powerpc64



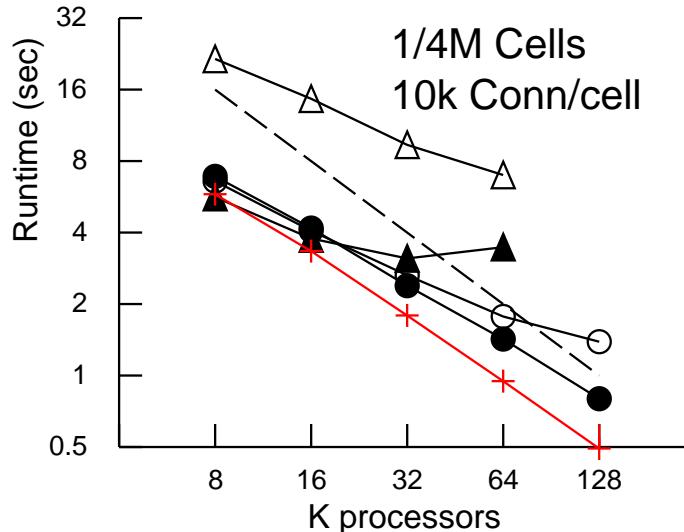
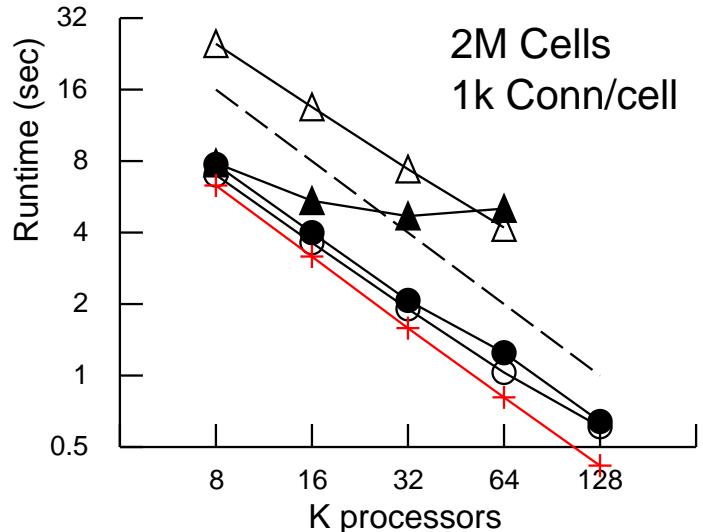
- \triangle MPI_ISend – Two Phase, Two Subinterval
- \blacktriangle Allgather
- \bullet DCMF_Multicast – Two Phase, Two Subinterval
- \circ Record-Replay – One Subinterval
- $+$ Computation Time (includes queue)

Artificial Spiking Net

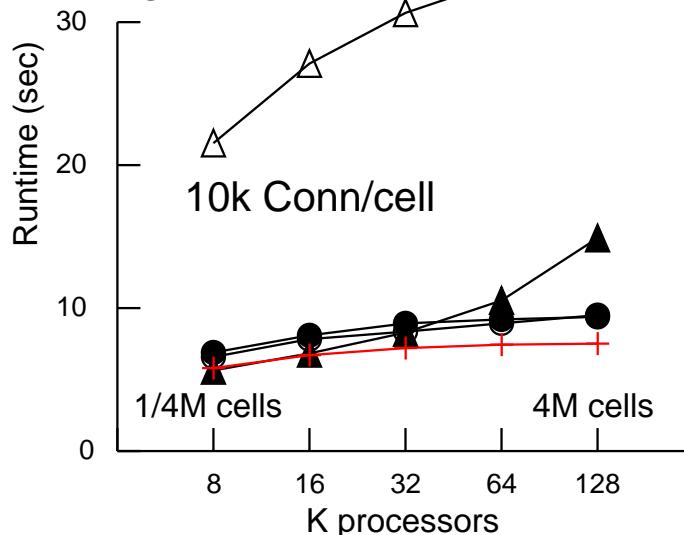
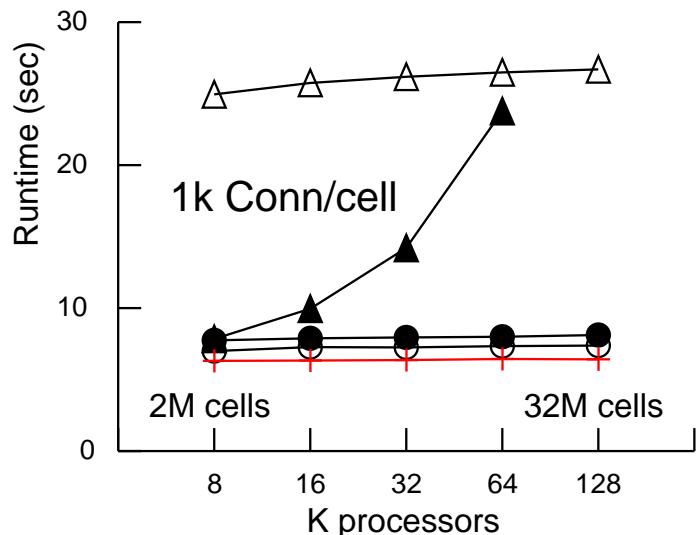
Blue Gene/P

Argonne National Lab

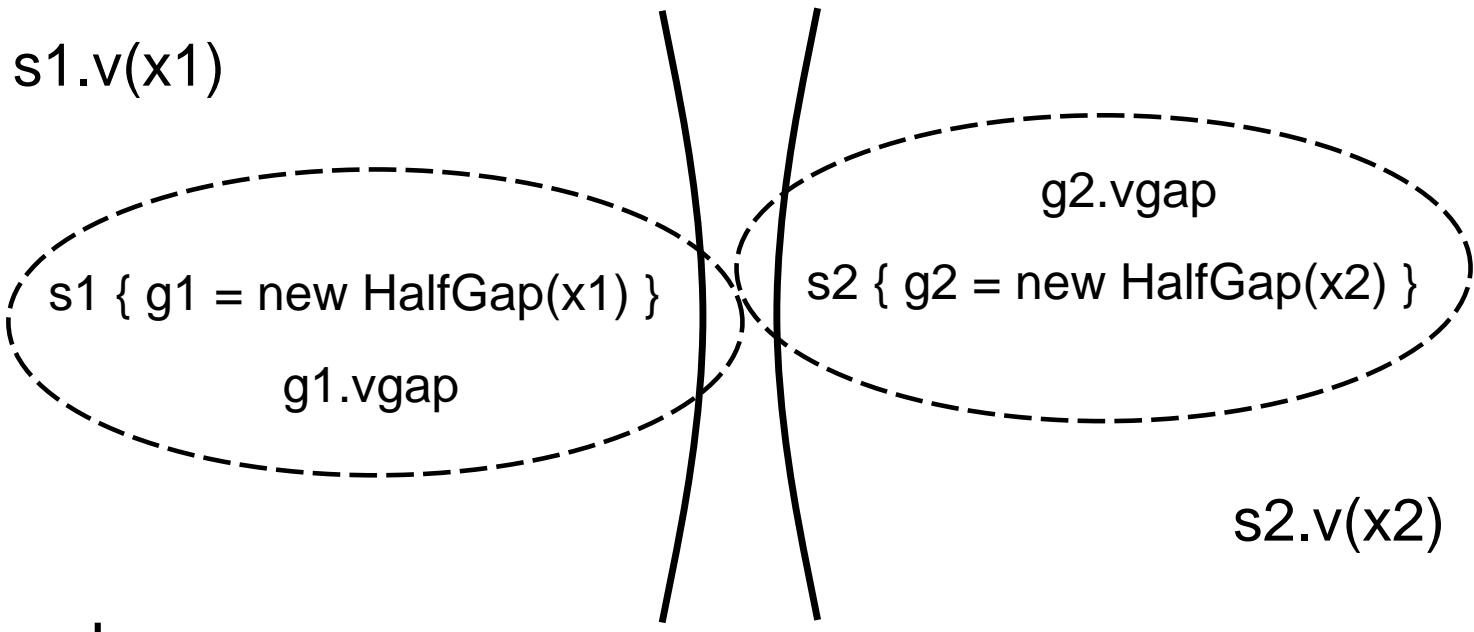
Strong Scaling



Weak Scaling



Continuous Voltage Exchange



gap.mod

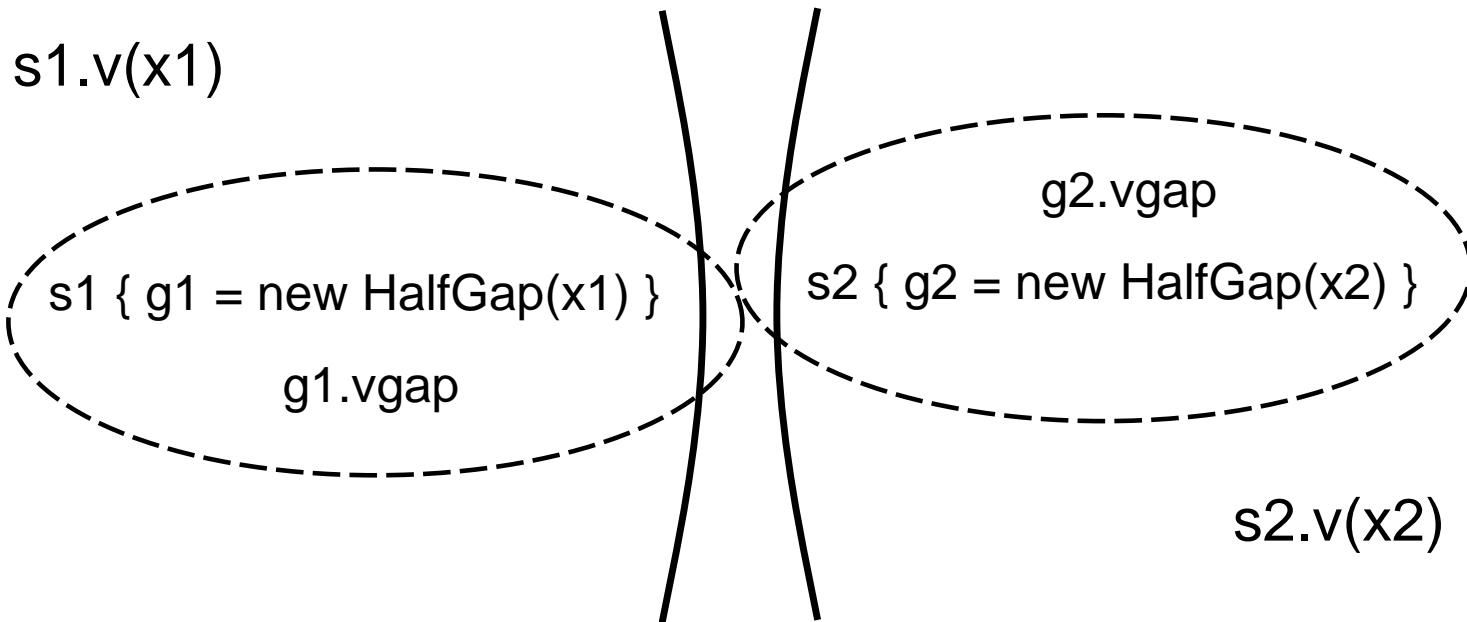
```
NEURON {  
    POINT_PROCESS HalfGap  
    ELECTRODE_CURRENT i  
    RANGE r, i, vgap  
}
```

```
PARAMETER { r = 1e9 (megohm) }
```

```
ASSIGNED {  
    v (millivolt)  
    vgap (millivolt)  
    i (nanoamp)  
}  
CURRENT { i = (vgap - v)/r }
```

Continuous Voltage Exchange

pc.source_var(&source_var, sgid)



gap.mod

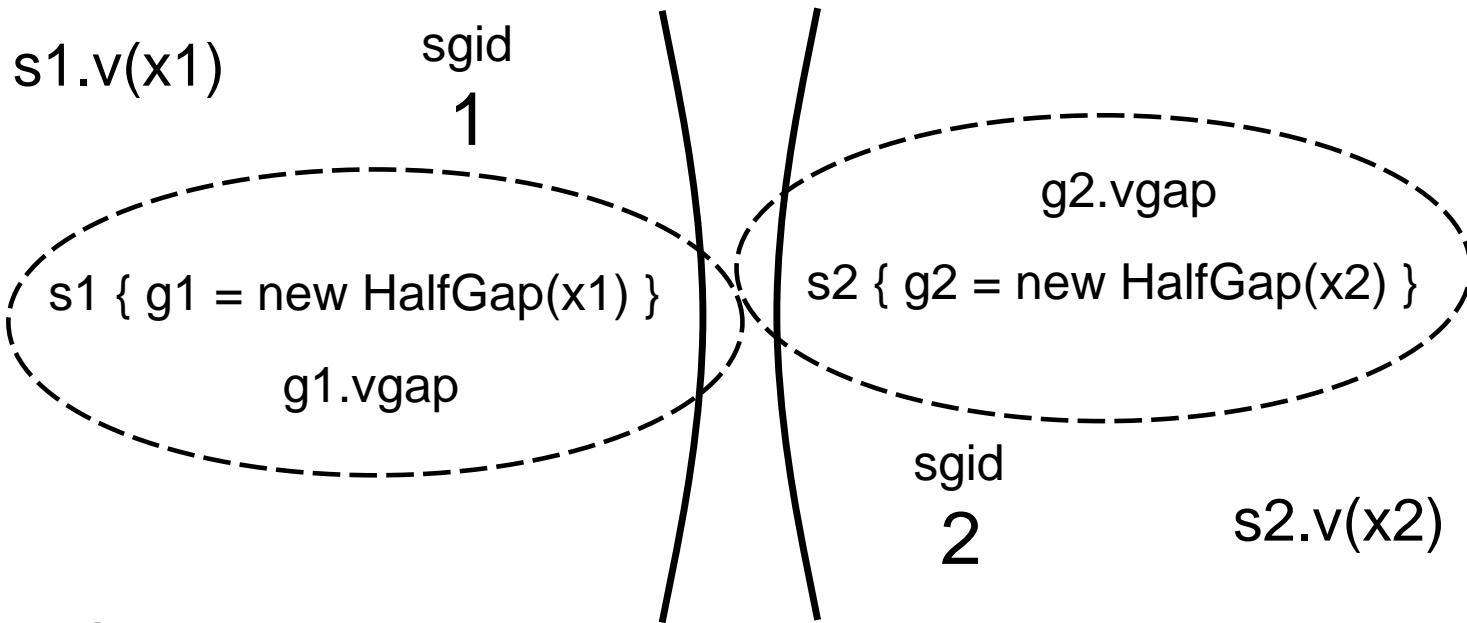
```
NEURON {
  POINT_PROCESS HalfGap
  ELECTRODE_CURRENT i
  RANGE r, i, vgap
}
```

```
PARAMETER { r = 1e9 (megohm) }
```

```
ASSIGNED {
  v (millivolt)
  vgap (millivolt)
  i (nanoamp)
}
CURRENT { i = (vgap - v)/r }
```

Continuous Voltage Exchange

pc.source_var(&source_var, sgid)



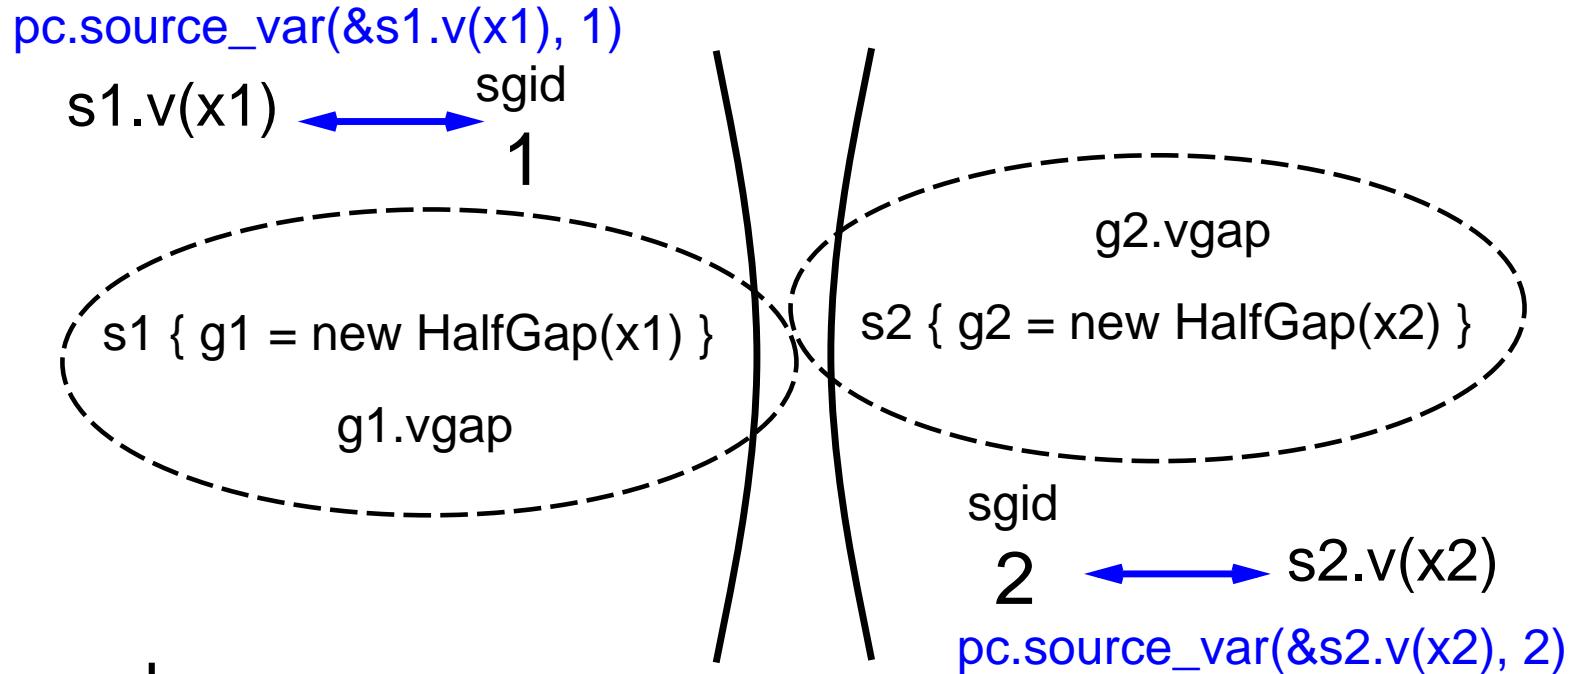
gap.mod

```
NEURON {  
    POINT_PROCESS HalfGap  
    ELECTRODE_CURRENT i  
    RANGE r, i, vgap  
}  
PARAMETER { r = 1e9 (megohm) }
```

```
ASSIGNED {  
    v (millivolt)  
    vgap (millivolt)  
    i (nanoamp)  
}  
CURRENT { i = (vgap - v)/r }
```

Continuous Voltage Exchange

pc.source_var(&source_var, sgid)



`gap.mod`

```
NEURON {  
    POINT_PROCESS HalfGap  
    ELECTRODE_CURRENT i  
    RANGE r, i, vgap  
}  
PARAMETER { r = 1e9 (megohm) }
```

```
ASSIGNED {  
    v (millivolt)  
    vgap (millivolt)  
    i (nanoamp)  
}  
CURRENT { i = (vgap - v)/r }
```

Continuous Voltage Exchange

pc.source_var(&source_var, sgid)

pc.target_var(&target_var, sgid)

pc.source_var(&s1.v(x1), 1)

s1.v(x1) \longleftrightarrow sgid
1

s1 { g1 = new HalfGap(x1) }

g1.vgap

g2.vgap

s2 { g2 = new HalfGap(x2) }

sgid
2 \longleftrightarrow s2.v(x2)

pc.source_var(&s2.v(x2), 2)

gap.mod

```
NEURON {
  POINT_PROCESS HalfGap
  ELECTRODE_CURRENT i
  RANGE r, i, vgap
}
```

```
PARAMETER { r = 1e9 (megohm) }
```

```
ASSIGNED {
  v (millivolt)
  vgap (millivolt)
  i (nanoamp)
}
CURRENT { i = (vgap - v)/r }
```

Continuous Voltage Exchange

pc.source_var(&source_var, sgid)

pc.target_var(&target_var, sgid)

pc.source_var(&s1.v(x1), 1)

s1.v(x1) \longleftrightarrow sgid
1

pc.target_var(&g2.vgap, 1)

g2.vgap

s1 { g1 = new HalfGap(x1) }

g1.vgap

pc.target_var(&g1.vgap, 2)

sgid
2

\longleftrightarrow s2.v(x2)

pc.source_var(&s2.v(x2), 2)

gap.mod

```
NEURON {
  POINT_PROCESS HalfGap
  ELECTRODE_CURRENT i
  RANGE r, i, vgap
}
```

```
PARAMETER { r = 1e9 (megohm) }
```

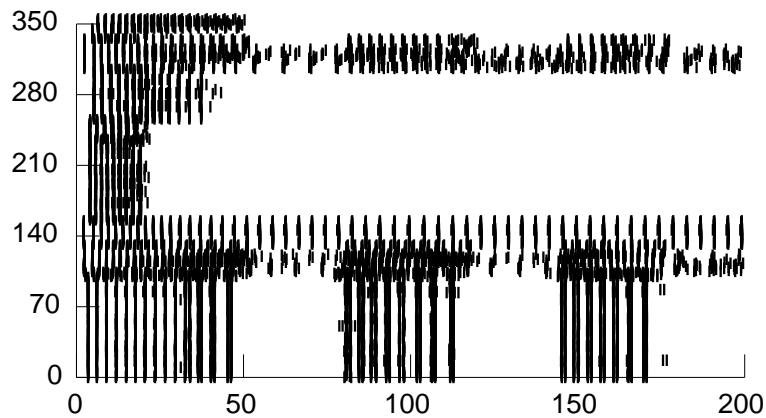
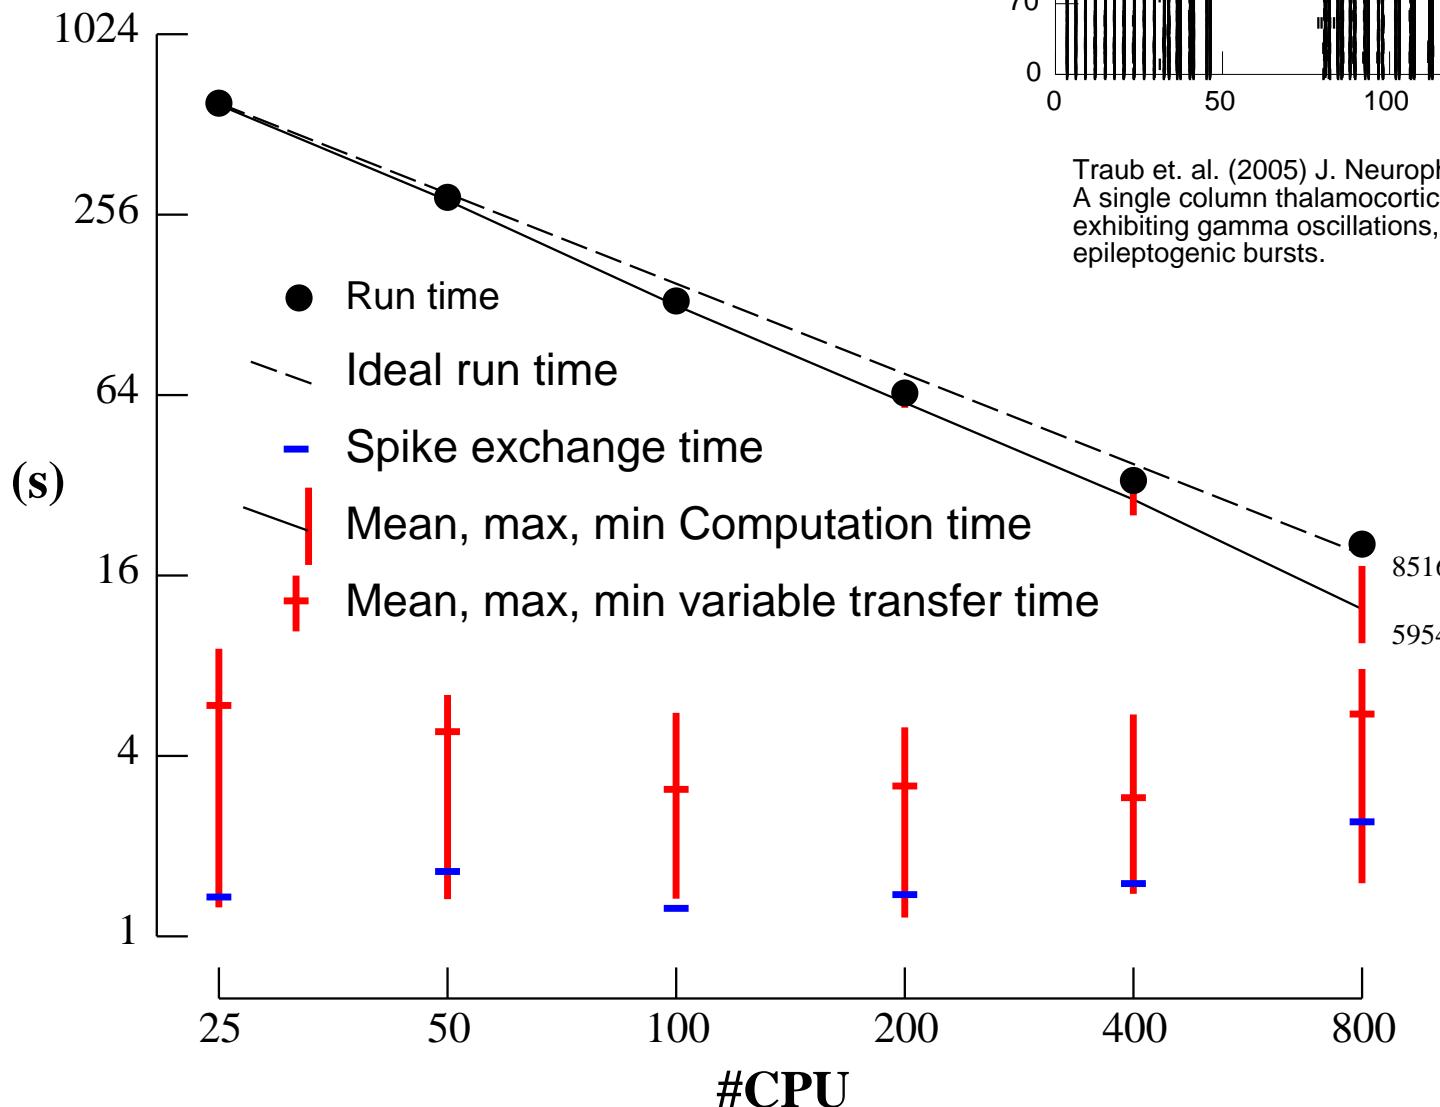
```
ASSIGNED {
  v (millivolt)
  vgap (millivolt)
  i (nanoamp)
}
CURRENT { i = (vgap - v)/r }
```

Pittsburgh Supercomputing Center

Bigben

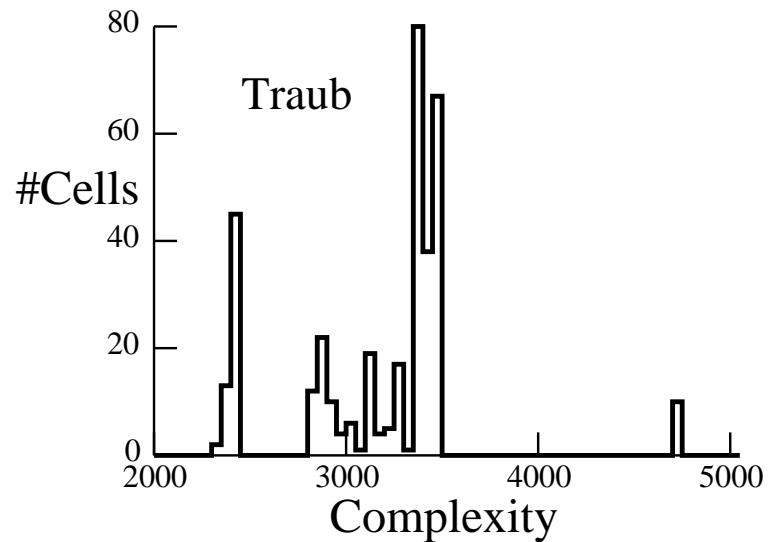
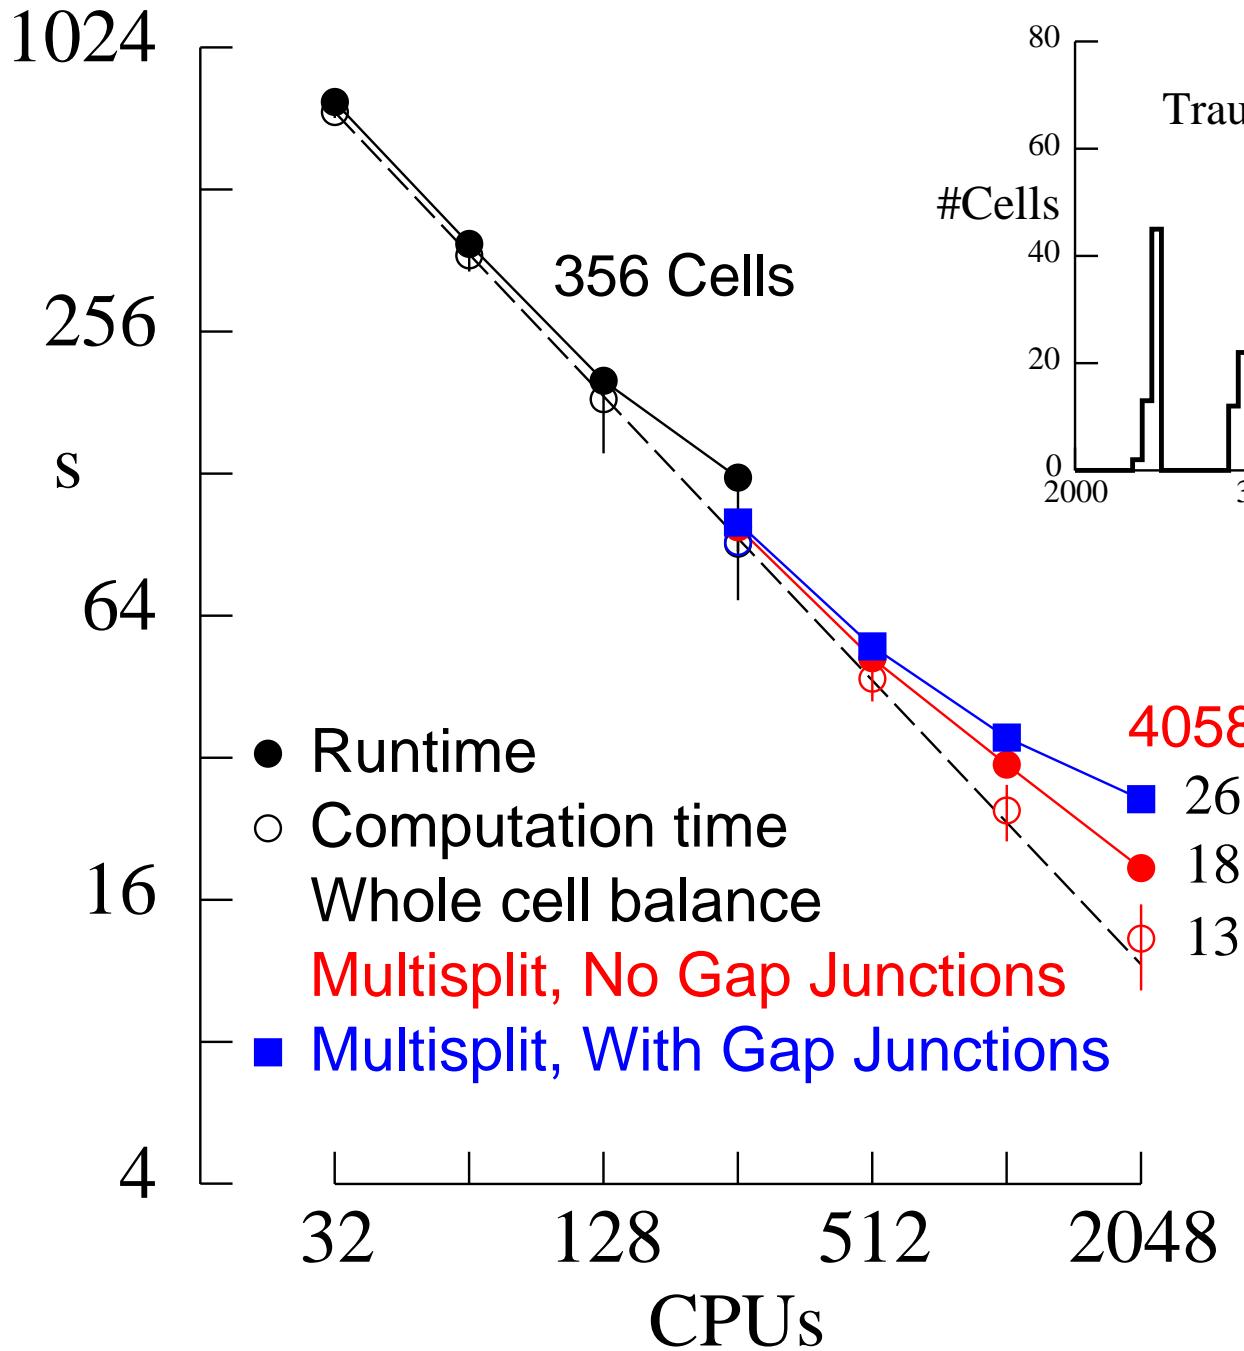
Cray XT3

2068 2.4 GHz Opteron Processors



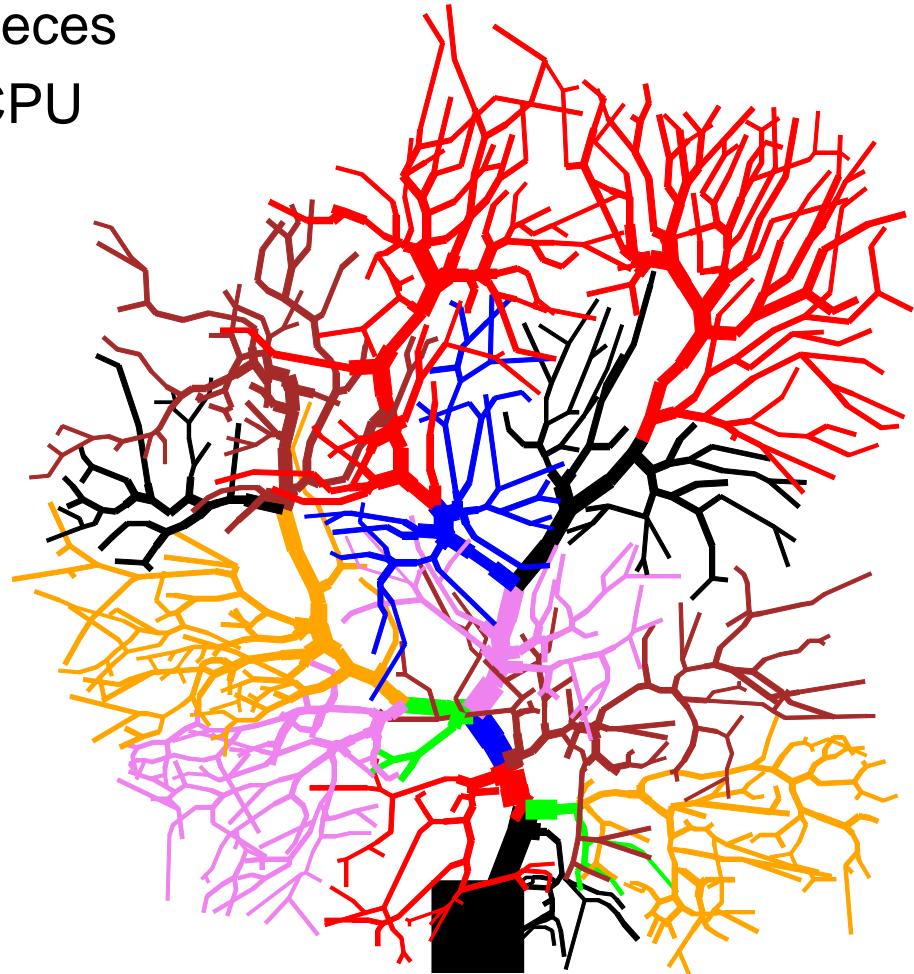
Traub et. al. (2005) J. Neurophysiol 93: 2194
A single column thalamocortical network model
exhibiting gamma oscillations, sleep spindles and
epileptogenic bursts.

3560 cells 14 types
3500 gap junctions
5,596,810 equations
73,465 spikes
1,122,520 connections
19,844,187 delivered



16 Pieces

4 CPU



| CPU | Time (s) | |
|-----|-------------|----------|
| | Computation | Exchange |
| 0 | 13.82 | 0.56 |
| 1 | 13.35 | 1.03 |
| 2 | 13.47 | 0.90 |
| 3 | 13.56 | 0.82 |

16 pieces, 1 cpu
wholecell, 1 cpu
16 pieces, 4 cpu

Runtime(s)
55.0
56.2
14.4

