

# The PoochMPI Toolkit

## High-Performance Computing with *Mathematica*

Wolfram Technology Conference 2006  
Champaign, Illinois

# **The PoochMPI Toolkit**

## **High-Performance Computing with *Mathematica***

**Dean E. Dager, Ph. D.**  
President, Dager Research, Inc.  
[d@daugerresearch.com](mailto:d@daugerresearch.com)

# The First Mac Cluster - established 1998



- 8 G3/266's
- Mac OS 8.1



- Decyk, Kokelaar, Dauger



# The First Mac Cluster - established 1998



4 Java Applets  
Distributed learning and teaching tools at California State University's Biology Life On-Line.

6 Technology Spotlight  
New tools and technologies to integrate Macintosh computing and the Internet.

6 Ask Apple  
Mac OS Runtime for Java, web browsers, and Java development tools for Macintosh.

7 Punching in  
At Cornell University, a Java application puts the time card process for 20,000 employees online.

www.apple.com News for the Academic Community Vol. 1 No. 1 Winter 1998

## Apple University Arts

**Profile:**  
UCLA's Project Appleseed  
(supercomputing for the rest of us).



### Project Appleseed

A UCLA physics lab uses a Power Macintosh cluster to run parallel processing super-computer software.

With software created by Decyk, Kokelaar, and Dauger at UCLA's Plasma Physics Center, the cluster of Mac computers is robust and very fast.

Items have come in recent years. For someone like me, who used the G4 VMP, this is shocking!"

Mac OS, compile your code with MacMP, and run! Network creation will be even easier with a cluster of Mac computers, since Decyk, Kokelaar, and Dauger built it.



- Decyk, Kokelaar, Dauger



# The First Mac Cluster - established 1998



**Profile:**  
UCLA's Project Appleseed  
(supercomputing for the rest of us).



## The Dawson Cluster





## High-Performance, Scientific, and Cluster Computing

- Software
  - “Plug-and-Play” Supercomputer-Compatible Clusters
  - Pooch Application
  - Source-Code Tutorials
  - Visualization & Simulation
- Consulting Services
  - Optimization
  - Parallelization
  - Vectorization

## High-Performance, Scientific, and Cluster Computing

- Software
  - “Plug-and-Play” Supercomputer-Compatible Clusters
  - Pooch Application
  - Source-Code Tutorials
  - Visualization & Simulation
- Consulting Services
  - Optimization
  - Parallelization
  - Vectorization



# Why Parallel Computing?





# Why Parallel Computing?

## Problems too large to solve on one computer


- Takes too much time
- Requires more memory
  - can outgrow RAM capacity

## Programming API standardized

- Message-Passing Interface (MPI)
  - specification established in 1994
  - dominant software interface at supercomputing centers
  - portable MPI code in Fortran and C




# MPI in *Mathematica*





# MPI in *Mathematica*



## PoochMPI Toolkit for *Mathematica*

- Closely follows MPI standard in *Mathematica* environment
  - Basic MPI calls (`mpiSend`, `mpiRecv`)
  - Asynchronous MPI calls (`mpiSend`, `mpiRecv`, `mpiTest`)
  - Collective MPI calls (`mpiBcast`, `mpiGather`, `mpiAlltoall`)
- High-level parallel calls for common tasks
  - `ParallelTable`, `EdgeCell`, `ParallelFourier`, `ElementManage`
- Basic Parallel I/O
- Automatically locates, launches, configures, and coordinates *Mathematica* kernels via Pooch
  - from command line or *Mathematica*'s Front End
- Builds on any licensed *Mathematica* (*LM*, *net-*, *grid-*, etc.)



# Why MPI instead of “grid”?





# Why MPI instead of “grid”?




# Why MPI instead of “grid”?






# Why MPI instead of “grid”?



ACS

# Why MPI instead of “grid”?


## Typical “grid” implementation





# Why MPI instead of “grid”?

## MPI implementation





# Parallel Code using MPI



# Parallel Code using MPI

## Code coordinating parallel work using messages

- N tasks or “virtual processors” running simultaneously, labeled 0 through N-1
- executables often use identification data to determine algorithmically what part of the problem on which to work
- tasks pass messages amongst themselves to organize data and coordinate work
- any group of tasks can communicate ( $\Rightarrow O(N^2)$  connections)
  - simple sends and receives
  - collective calls - broadcast, gather, reduce, transpose, etc.
- synchronization not required, but often implied by messages



# “Game of Life”

## Cellular Automata

- J. Conway, Princeton

```
030      10      02      040 3 0 2      3      2
      21      14      2      21      0      1
020      0      02 021      1 1
0 3      110      0      3 30      75 0
2      5      1 1      4
0 30 010      0      0 * 0      4
05 2 022      1 1
      00 0      01      74      020      1 0
      2      320      3 3      0 1
*      2      02 1      45      10
* *      079      10
* *      59
*      64      0 40      2
      02 00      0
      6041      1 0
2 2002      0 2
      0 0      10
0*0 010 54      010
030      00      01
      0      13      010
0      056000 1 2      09*2010
```

Life of one cell depends on its neighbors





# “Game of Life”

## Cellular Automata

- J. Conway, Princeton

3	0	
	2	
		0

```

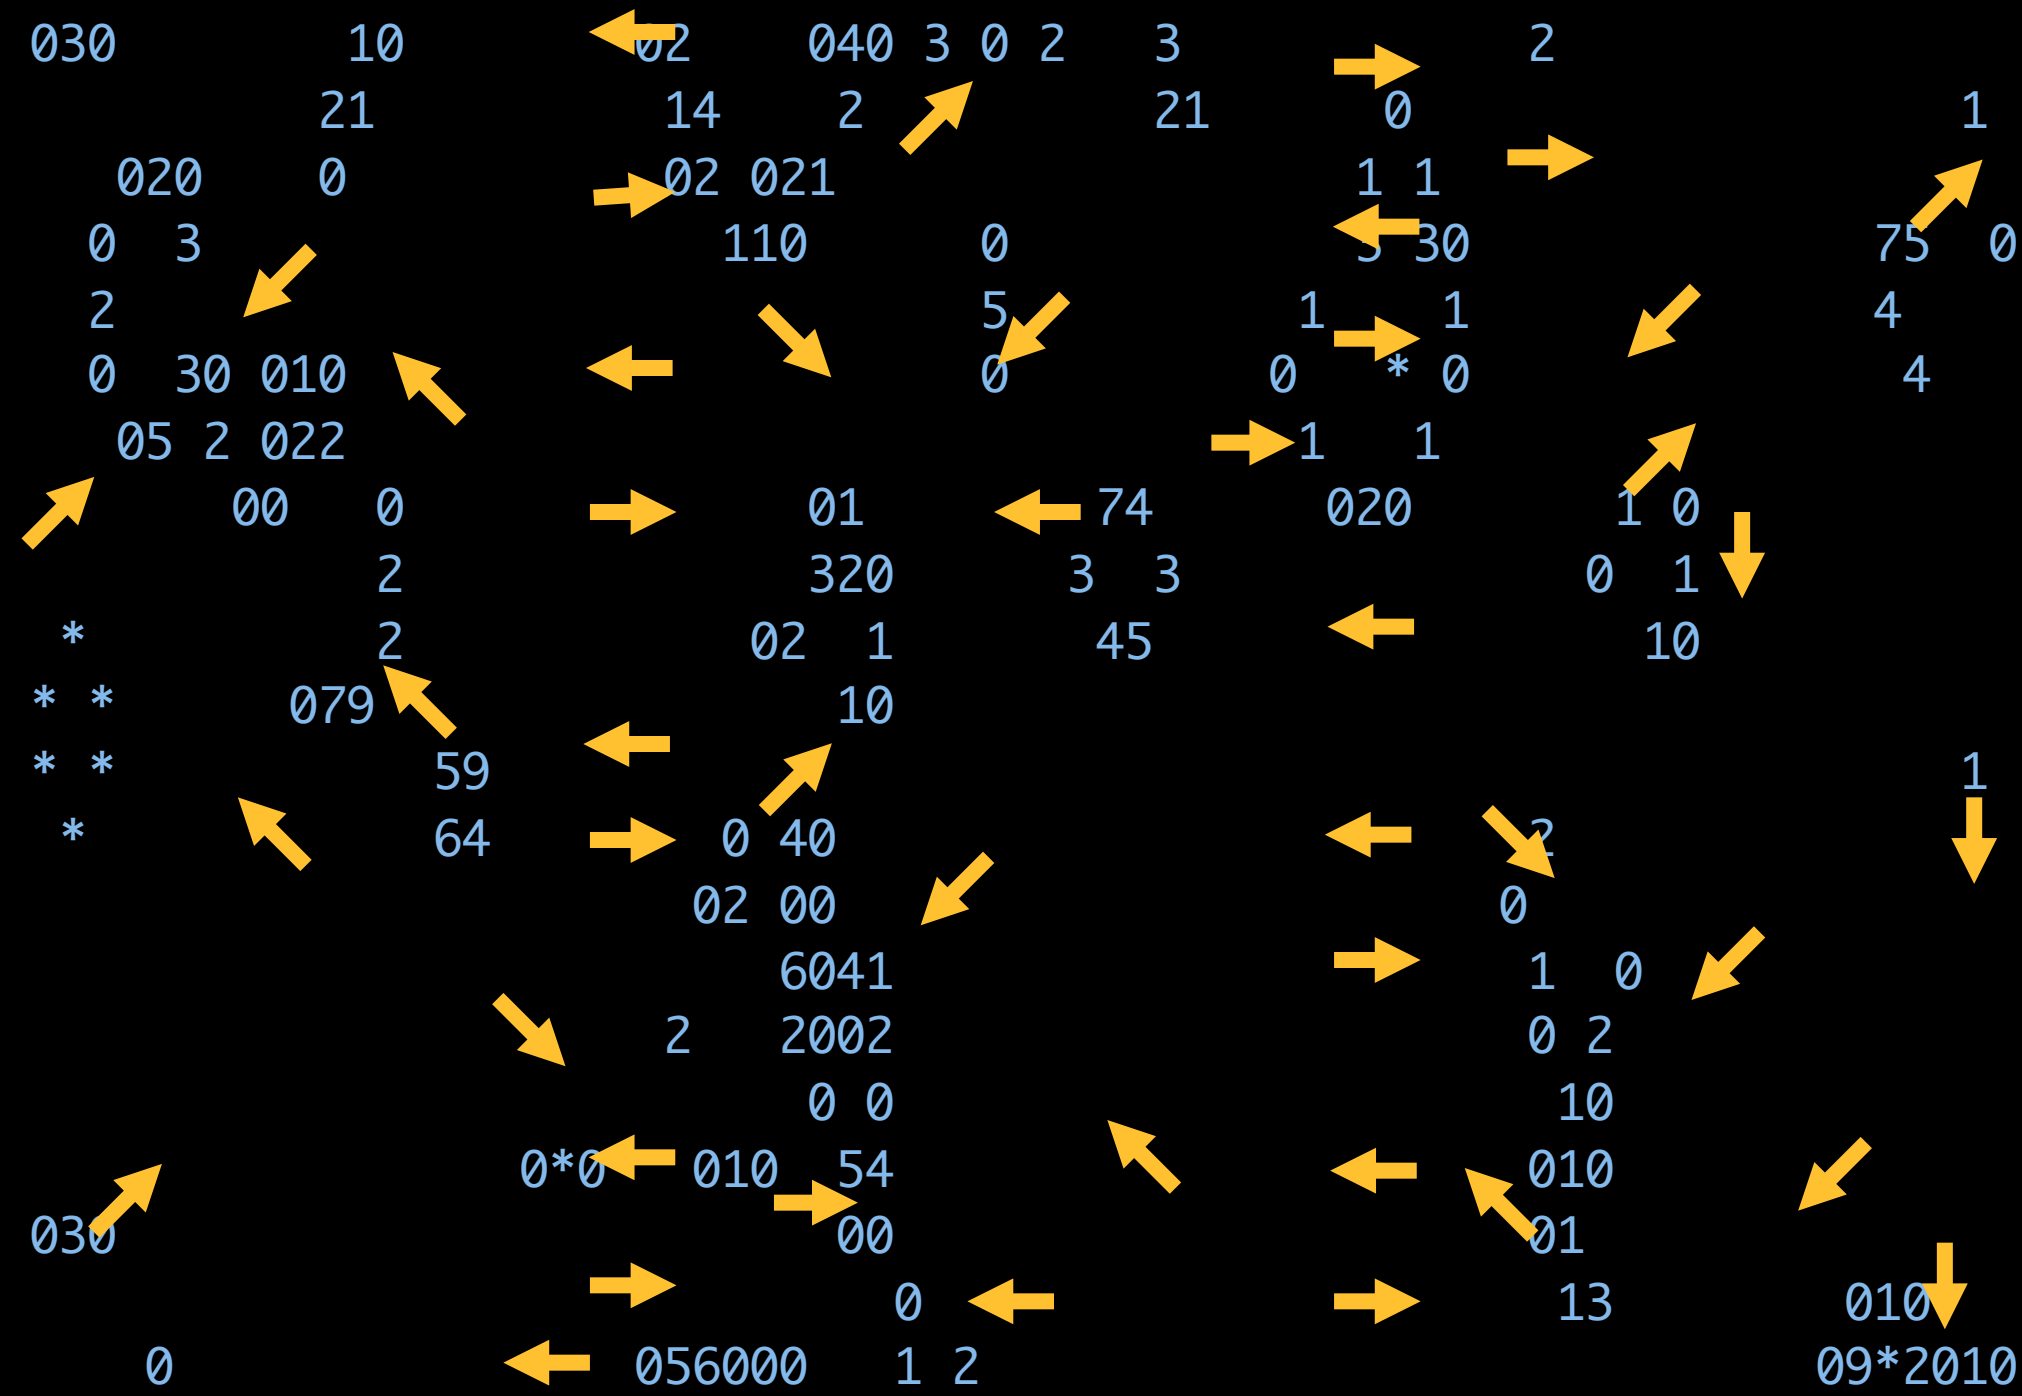

030      10      02      040 3 0 2      3      2
      21      14      2      21      0      1
020      0      02 021      1 1
0 3      110      0      3 30      75 0
2      5      1 1      4
0 30 010      0      0 * 0      4
05 2 022      1 1
      00      0      01      74      020      1 0
      2      320      3 3      0 1
*      2      02 1      45      10
* *      079      10
* *      59
*      64      0 40      2
      02 00      0
      6041      1 0
2 2002      0 2
      0 0      10
0*0 010 54      010
030      00      01
      0      13      010
0      056000 1 2      09*2010
    
```

Life of one cell depends on its neighbors



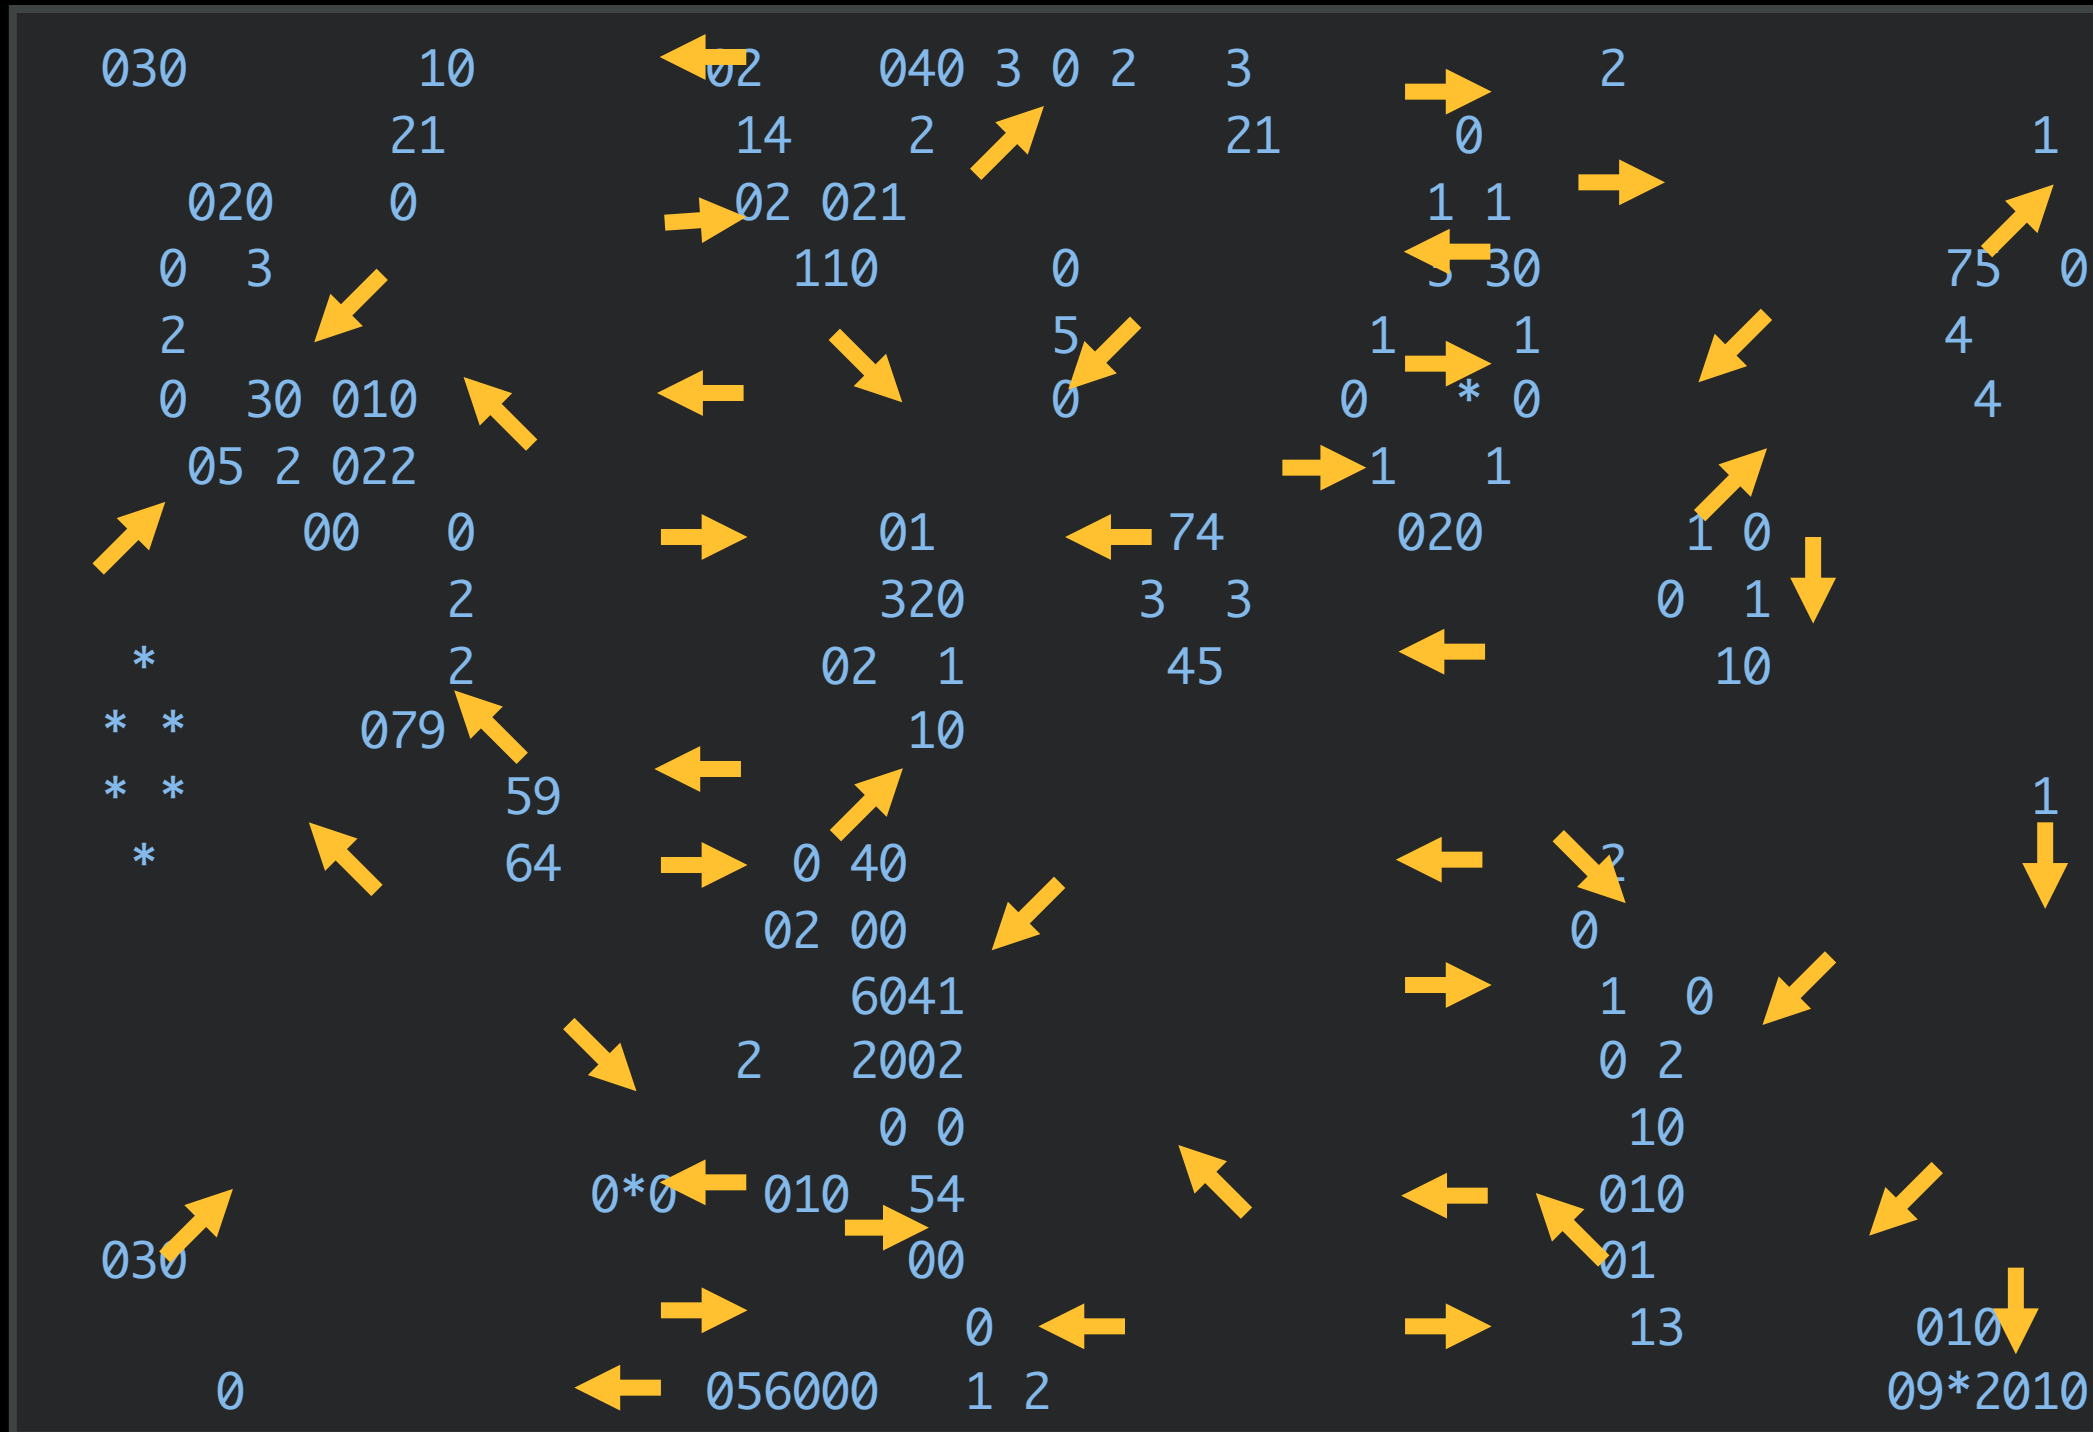

# Parallel "Life"

## Message-Passing for Cellular Automata



# Parallel "Life"



## Message-Passing for Cellular Automata





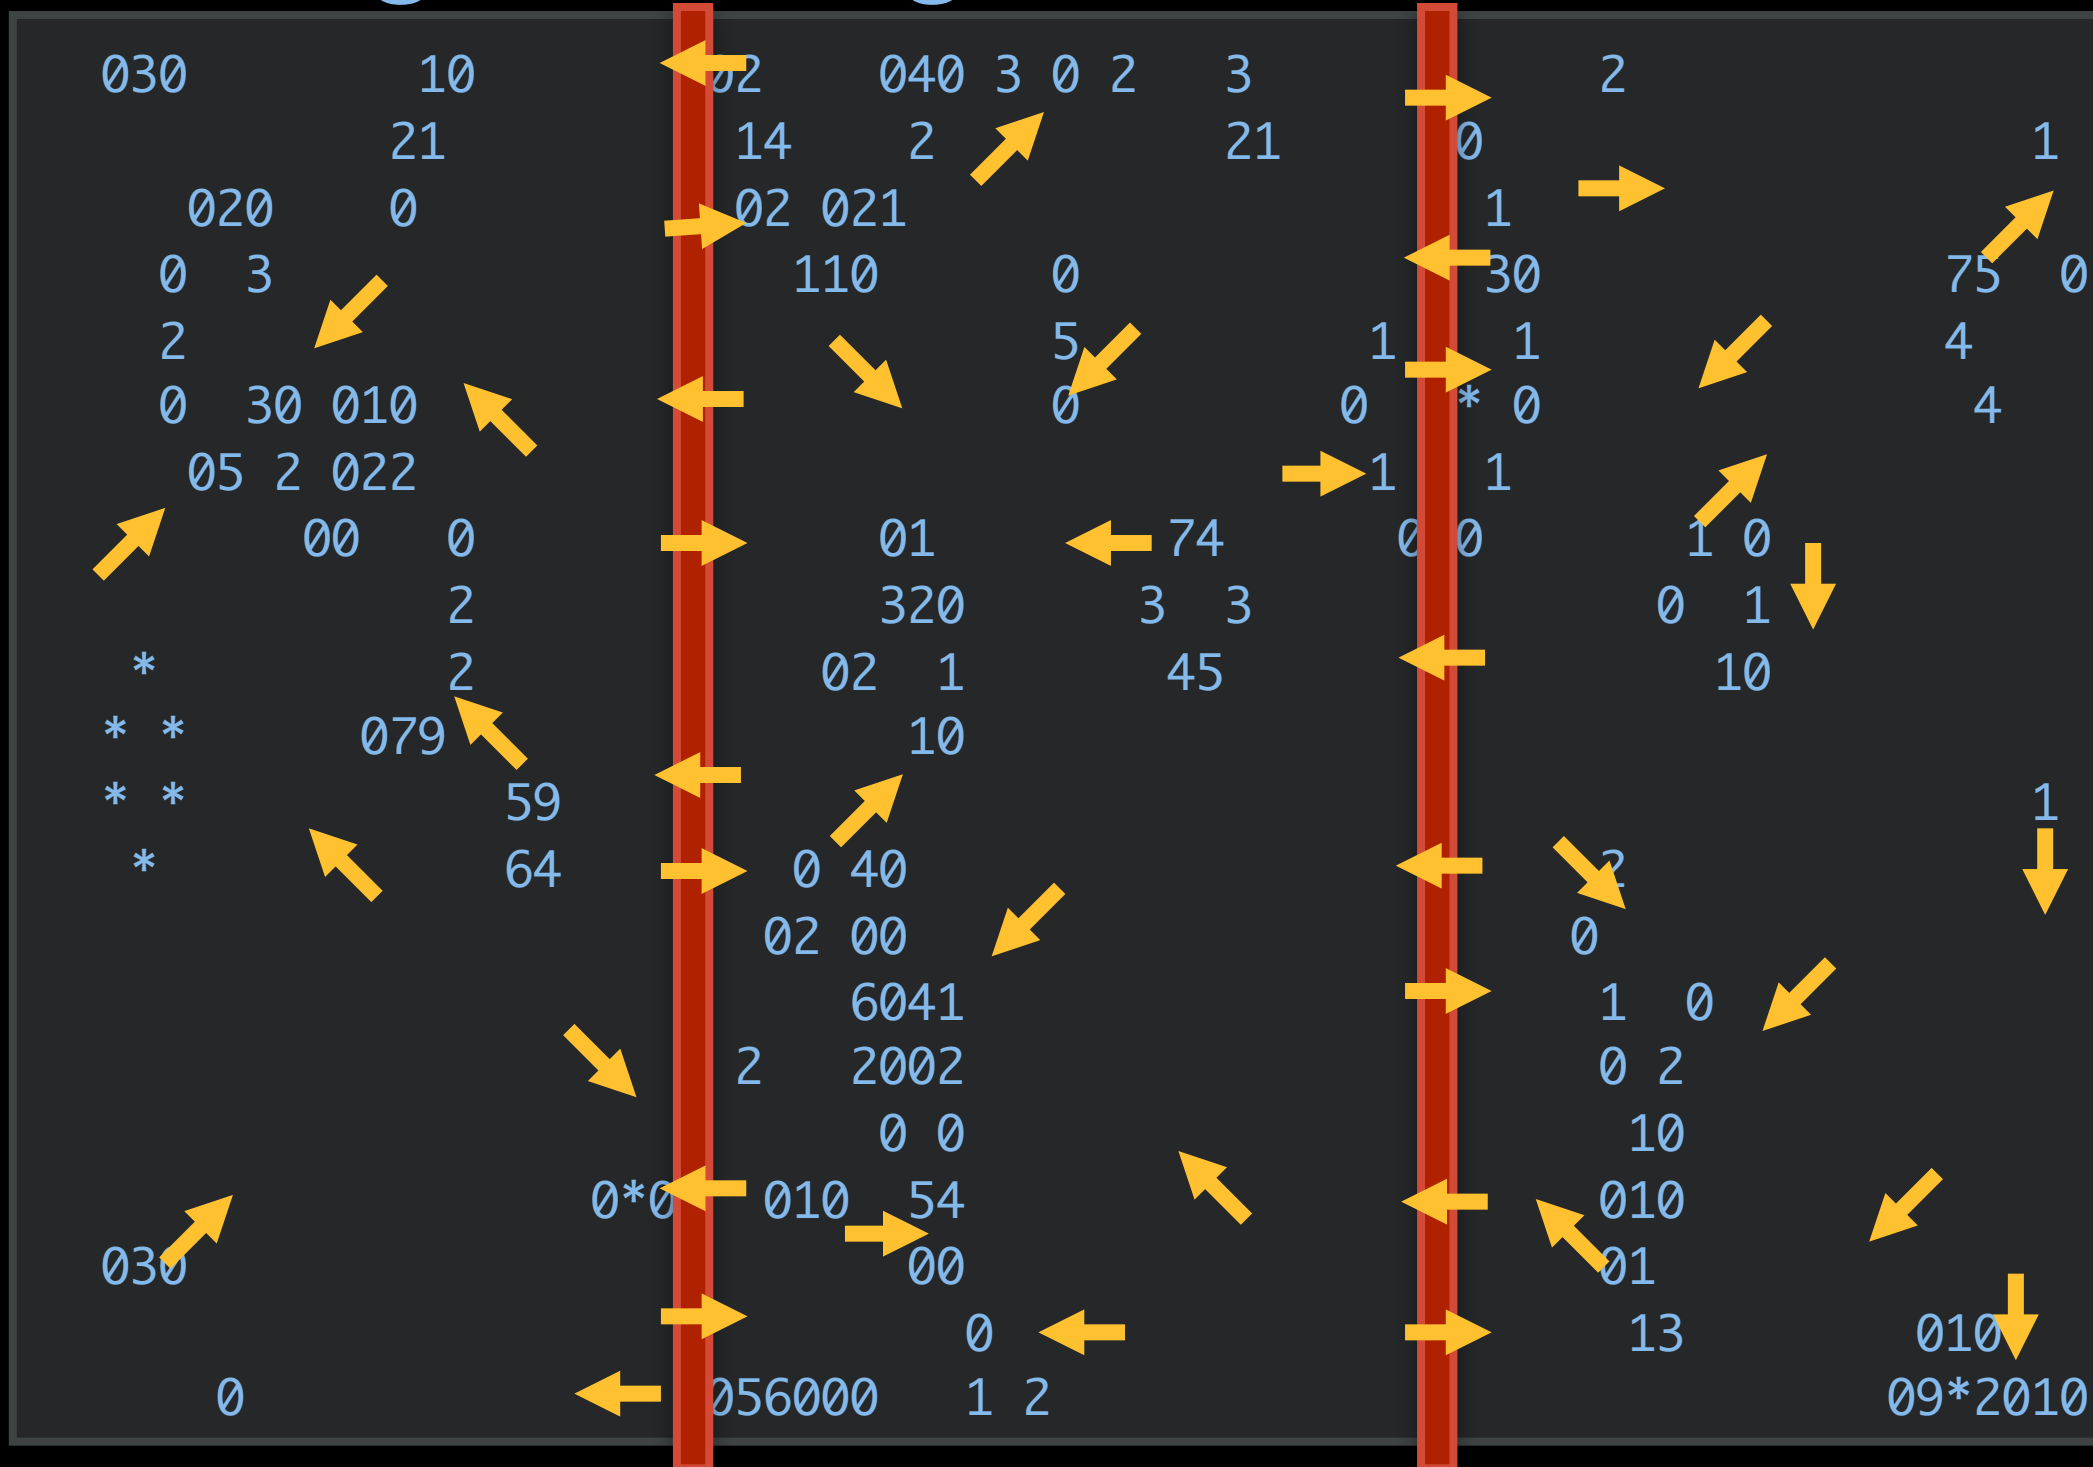

# Parallel "Life"

## Message-Passing for Cellular Automata



# Parallel "Life"

## Message-Passing for Cellular Automata



Data to propagate between partitions



# Parallel "Life"

## Message-Passing for Cellular Automata

```
030      10      0
          21
        020      0
        0 3
        2
        0 30 010
        05 2 022
          00      0
          2
        *      2
        * *      079
        * *      59
        *      64
          0*0
030
          0      0
```

```
2  040 3 0 2  3
4  2      21
02 021
  110      0
          5      1
          0      0      1
          1      0
          01      74      0
          320      3 3
        02 1      45
          10
          0 40
        02 00
          6041
        2  2002
          0 0
        010 54
          00
          0
        56000 1 2
```

```
          2
          0      1
        1 1
        3 30      75 0
          1      4
        * 0      4
          1
        20      1 0
          0 1
          10
          1
          2
          0
          1 0
          0 2
          10
          010
          01
          13      010
          09*2010
```



# Parallel "Life"

## Message-Passing for Cellular Automata

```
030      10      0
          21
        020      0
        0 3
        2
        0 30 010
        05 2 022
          00      0
            2
        *      2
        * *
        * *      079
        *      59
          64
          0*0
030
0
```

```
2      040 3 0 2 3
4      2      21
02 021
      110      0
          5
          0
          0
          01      74
          320      3 3
          02 1      45
          10
          0 40
          02 00
          6041
        2 2002
          0 0
          010 54
          00
          0
        5000 1 2
```

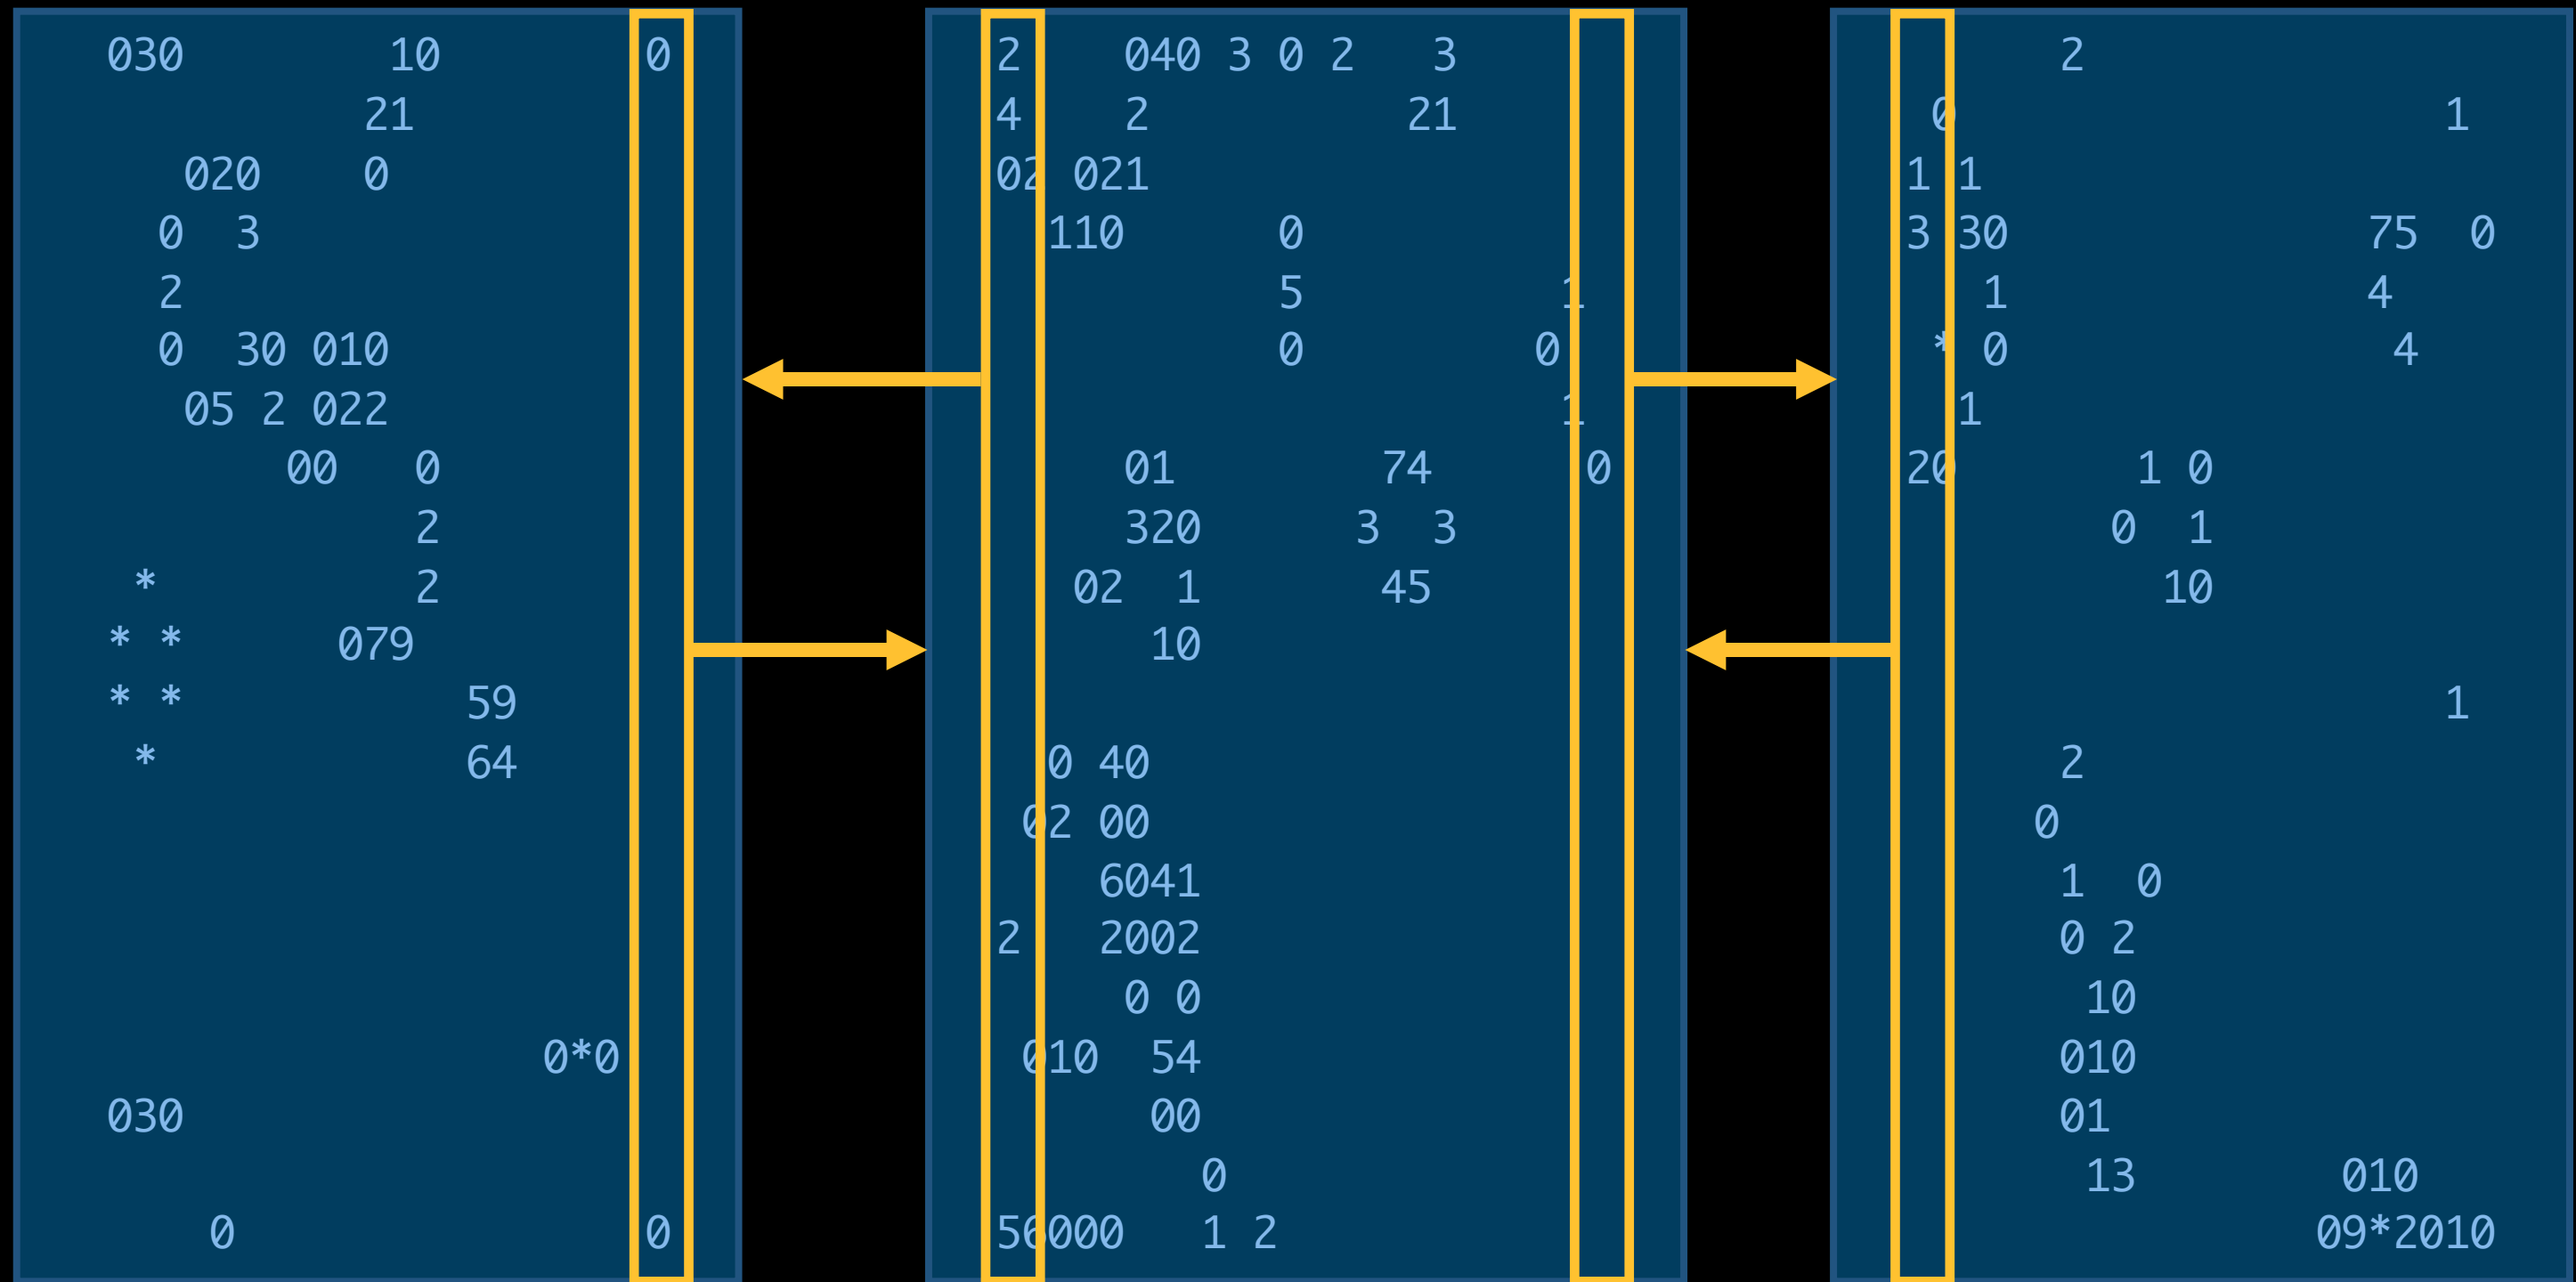
```
2
0      1
1 1
3 30      75 0
      1      4
      * 0      4
      1
      20      1 0
          0 1
          10
          1
          0 2
          10
          010
          01
          13      010
          09*2010
```






# Parallel "Life"

## Message-Passing for Cellular Automata



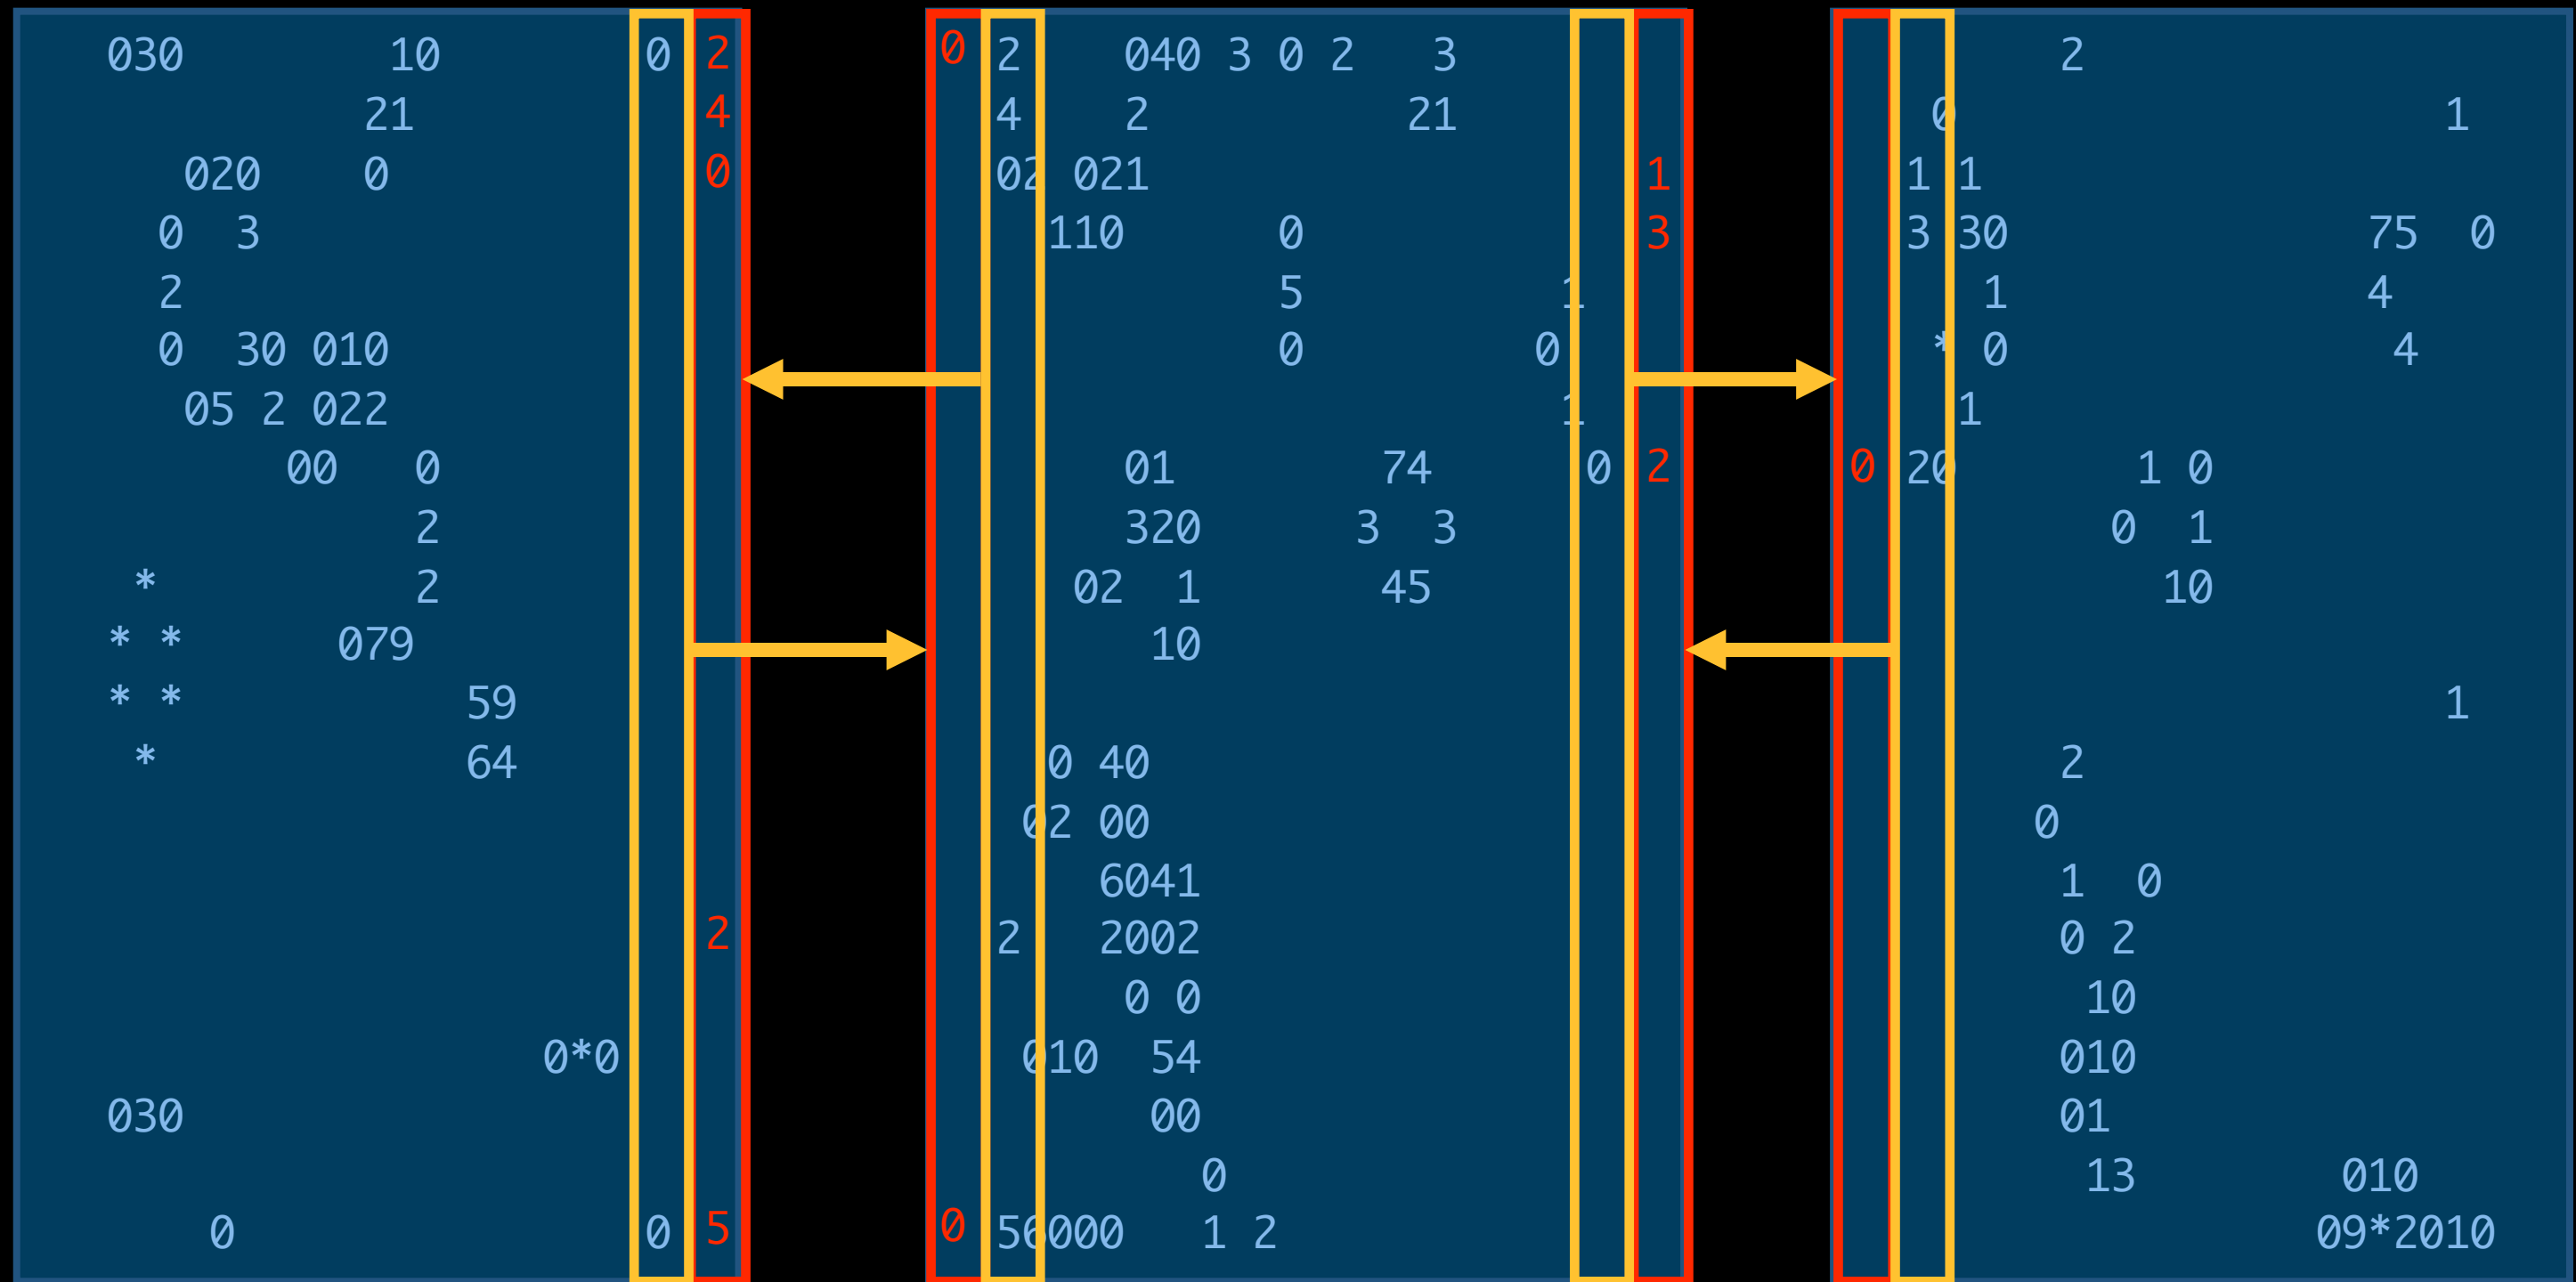
# Parallel "Life"

## Message-Passing for Cellular Automata



# Parallel "Life"

## Message-Passing for Cellular Automata




Message exchange maintains "guard cells"



# Plasma Simulation

## Plasma Dynamics



V. K. Decyk, C. D. Norton, *Comp. Phys. Communications* **164** (2004) 80-85





# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics


```
*      *      *      *      *      *      *      *      *      *      *      *      *      *      *
  ****  **      *      *      *      *      *      *      *      *      *      *      *      *
    *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
      *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
**   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
**   *   ****      *      *      *      *      *      *      *      *      *      *      *
*   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
    **   *      **   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *      *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
**   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  ****  *   *      *      *      *      *      *      *      *      *      *      *
    *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
*   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
    *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
**   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
    *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
*   **   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
  *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
    *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
```

V. K. Decyk, C. D. Norton, *Comp. Phys. Communications* **164** (2004) 80-85



# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics




V. K. Decyk, C. D. Norton, *Comp. Phys. Communications* **164** (2004) 80-85



# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics




V. K. Decyk, C. D. Norton, *Comp. Phys. Communications* **164** (2004) 80-85



# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics



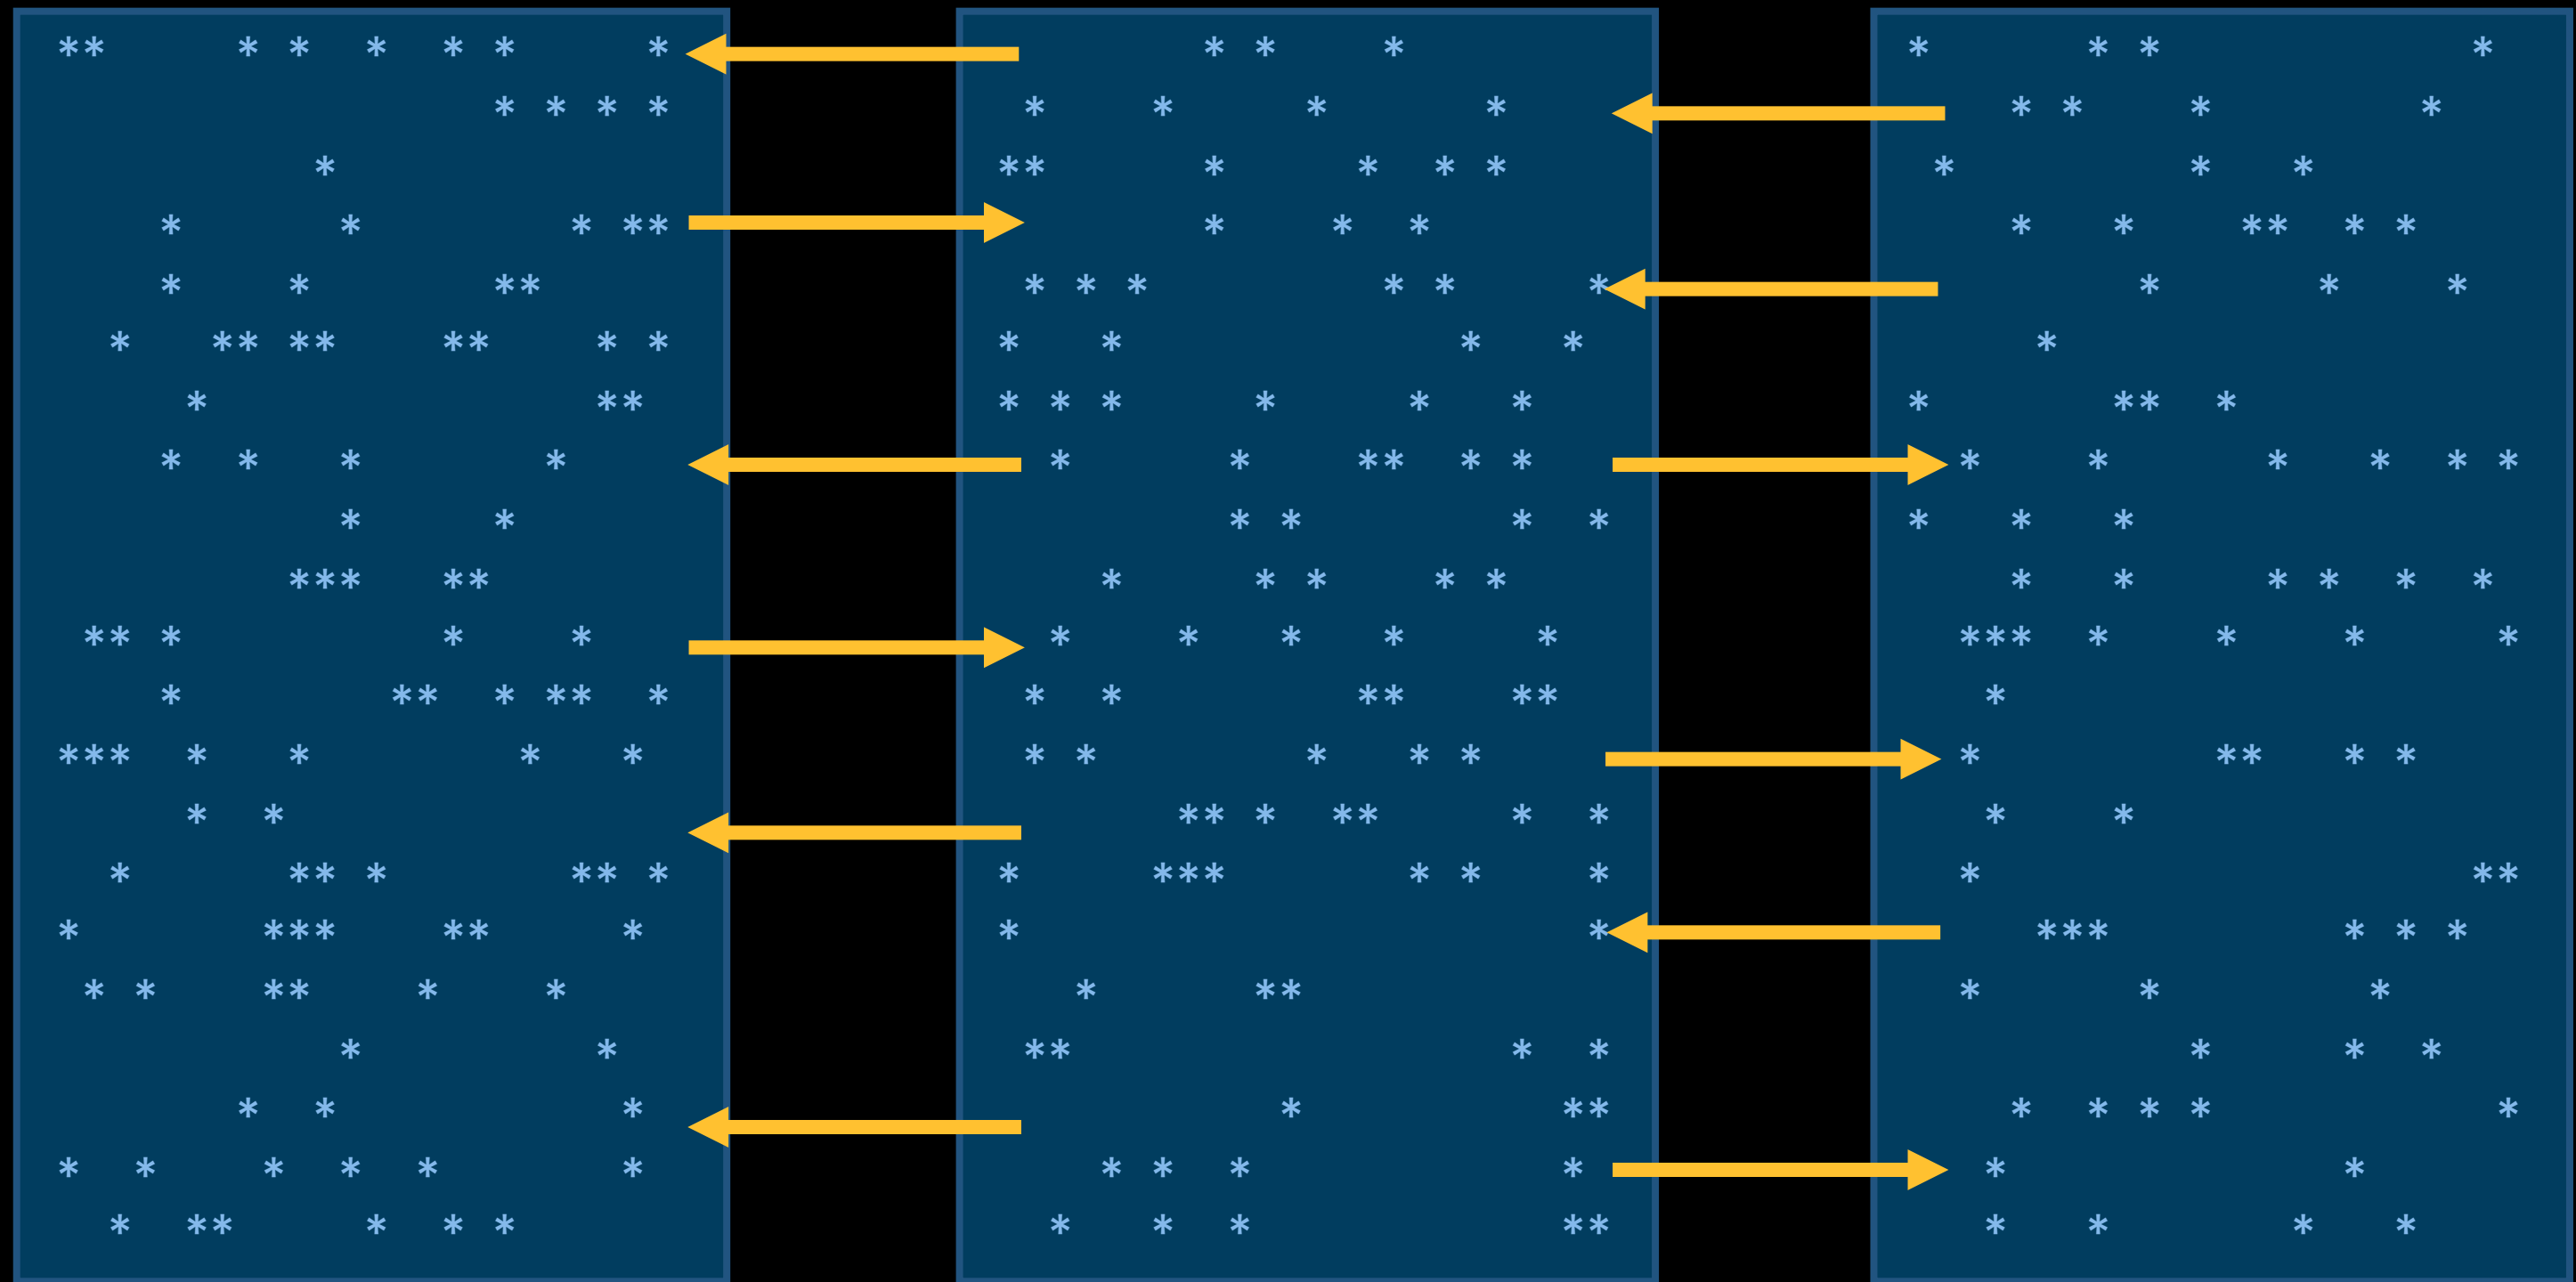
V. K. Decyk, C. D. Norton, *Comp. Phys. Communications* **164** (2004) 80-85





# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics




Particles propagate from one partition to its neighbor



# Parallel Plasma Simulation

## Message-Passing for Plasma Dynamics









Particles propagate from one partition to its neighbor




# Message-Passing Patterns

Supported via MPI




# Message-Passing Patterns

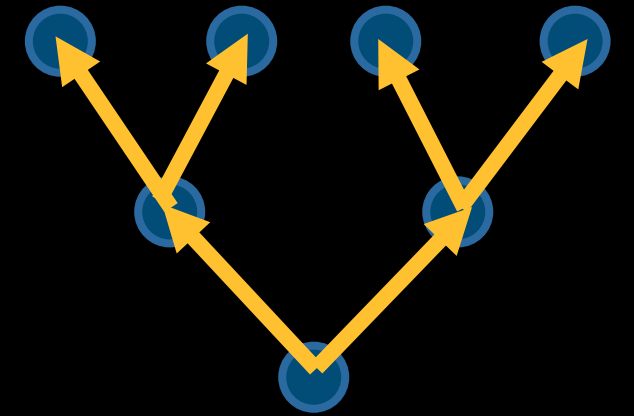
Supported via MPI




“Master-Slave”




All to All




Tree



Nearest Neighbor



Irregular



Or Combinations





# A Demonstration

Dean E. Dauger, Ph. D.  
d@daugerresearch.com



<http://daugerresearch.com>

# Who to Contact

---

**Dean E. Dauger, Ph. D.**

Dauger Research, Inc.

[d@daugerresearch.com](mailto:d@daugerresearch.com)

---

**Zvi Tannenbaum**

Advanced Cluster Systems, LLC

[zvi@advclustersys.com](mailto:zvi@advclustersys.com)

---

---

---



<http://daugerresearch.com> & <http://acs-grid.com>

**For More Information**

# Reference Library

## Documentation

### **PoochMPI for Mathematica Web Site**

<http://daugerresearch.com/pooch/mathematica/>

### **Advanced Cluster Systems Web Site**

<http://acs-grid.com/>

### **Tutorials on Writing Parallel Code**

<http://daugerresearch.com/pooch/tutorials.html>

### **Mac Clustering on National Television**

<http://daugerresearch.com/awards/KeepingAmericaStrong.html>

### **Related Publications**

<http://daugerresearch.com/pooch/publication.html>



# Reference Library

## Documentation

### **PoochMPI for Mathematica Web Site**

<http://daugerresearch.com/pooch/mathematica/>

### **Advanced Cluster Systems Web Site**

<http://acs-grid.com/>

### **Tutorials on Writing Parallel Code**


<http://daugerresearch.com/pooch/tutorials.html>

### **Mac Clustering on National Television**

<http://daugerresearch.com/awards/KeepingAmericaStrong.html>

## Related Publications

<http://daugerresearch.com/pooch/publication.html>





# Q&A

**Dean E. Dauger, Ph. D.**  
President, Dauger Research, Inc.  
d@daugerresearch.com



<http://daugerresearch.com> & <http://acs-grid.com/>