

# Towards parallelizing the Gillespie SSA

**Srivastav Ranganathan**  
and  
**Aparna JS**

**Indian Institute of Technology Bombay**  
**Mumbai,**  
**India**

# Gillespie Algorithm

- A stochastic simulation approach to study the time evolution of a system of reactions (processes)
- Each reaction occurs at an average rate
- The abundance of various species and their rates decide the propensity of each event to occur
- Many independent trajectories are generated to compute ensemble averaged statistical quantities

# Where is it used?

- In biological systems
- Outcomes of cellular processes are driven by stochasticity at a molecular level
- Deterministic approaches cannot capture the inherent randomness of biological systems

# The algorithm

- Assume a system whose in state A (its configuration)
- 'M' independent reactions are possible at any given time
- Each of these 'i' processes could occur with a rate,  $r_i$
- The probability of occurrence of each of these processes is thus given by

$$p_i = \frac{r_i}{\sum r_i}$$

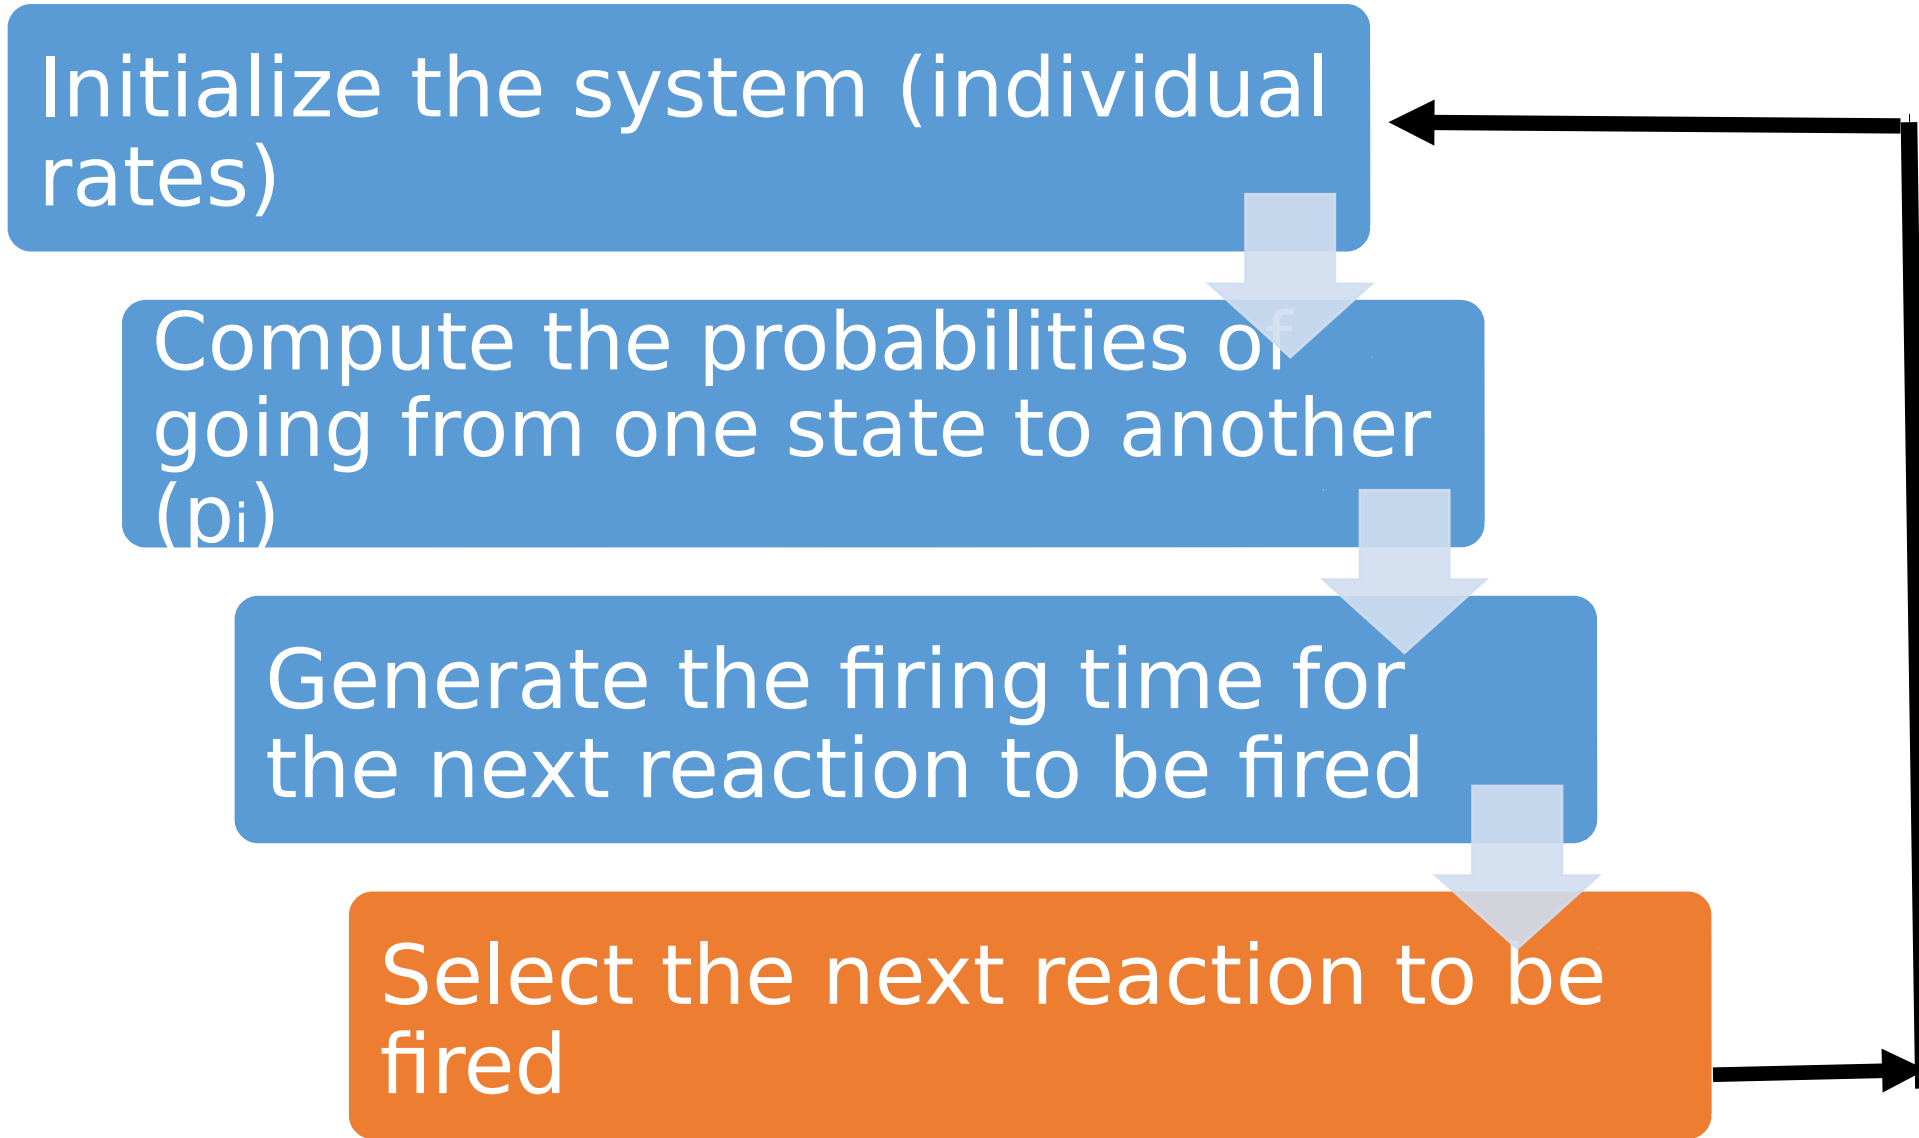
Initialize the system (individual rates)

Compute the probabilities of going from one state to another ( $p_i$ )

Generate the firing time for the next reaction to be fired

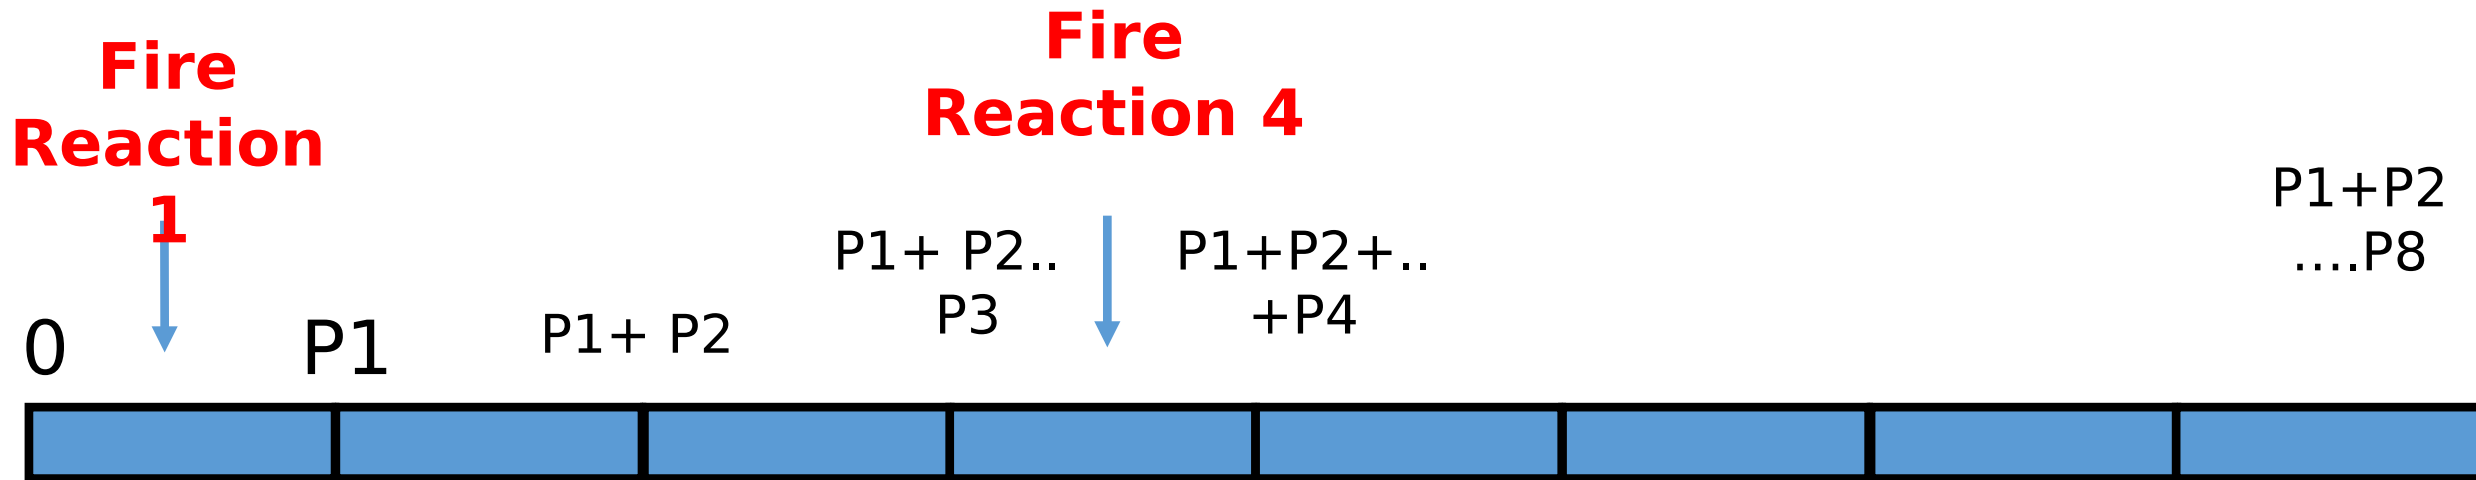
Select the next reaction to be fired

Most expensive of all these steps



# Selecting the event to be fired

- Draw a uniform random number **ran1**




- Update the system configuration based on the fired reaction (abundance, rates etc)
- Update the time based on the exponential distribution of wait time between events (

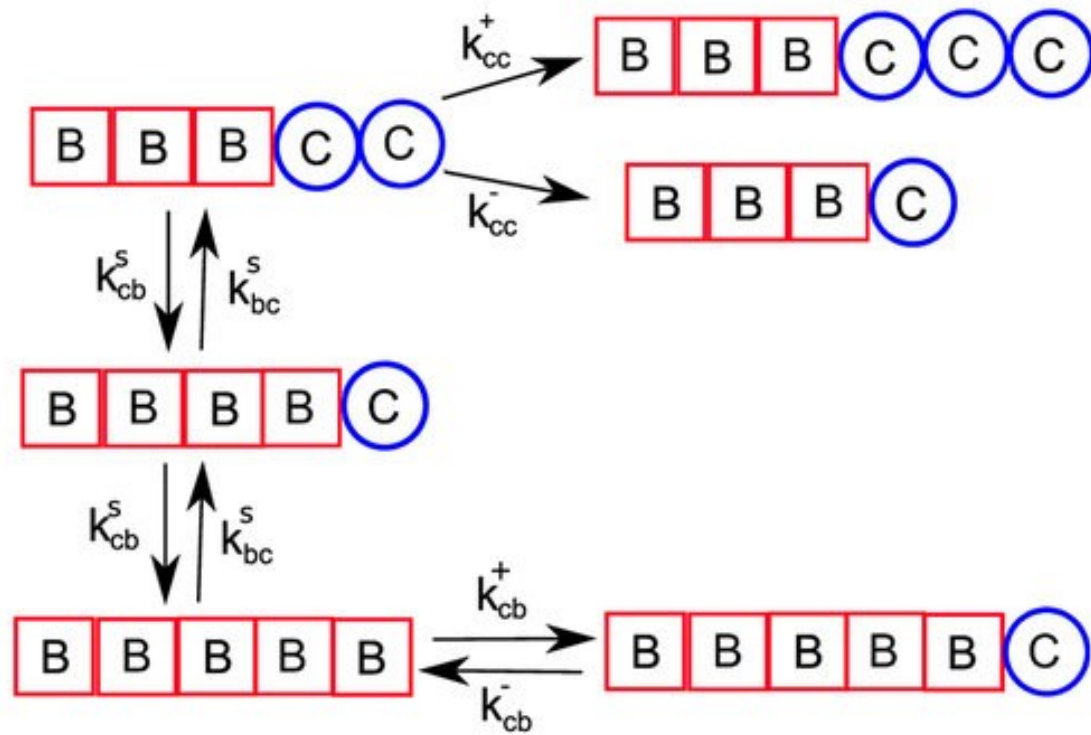
$$dt = \frac{1}{\sum r_i} \ln \frac{1}{ran2}$$

This search for the next event to be fired is really expensive if there is a large reaction space!

A



B

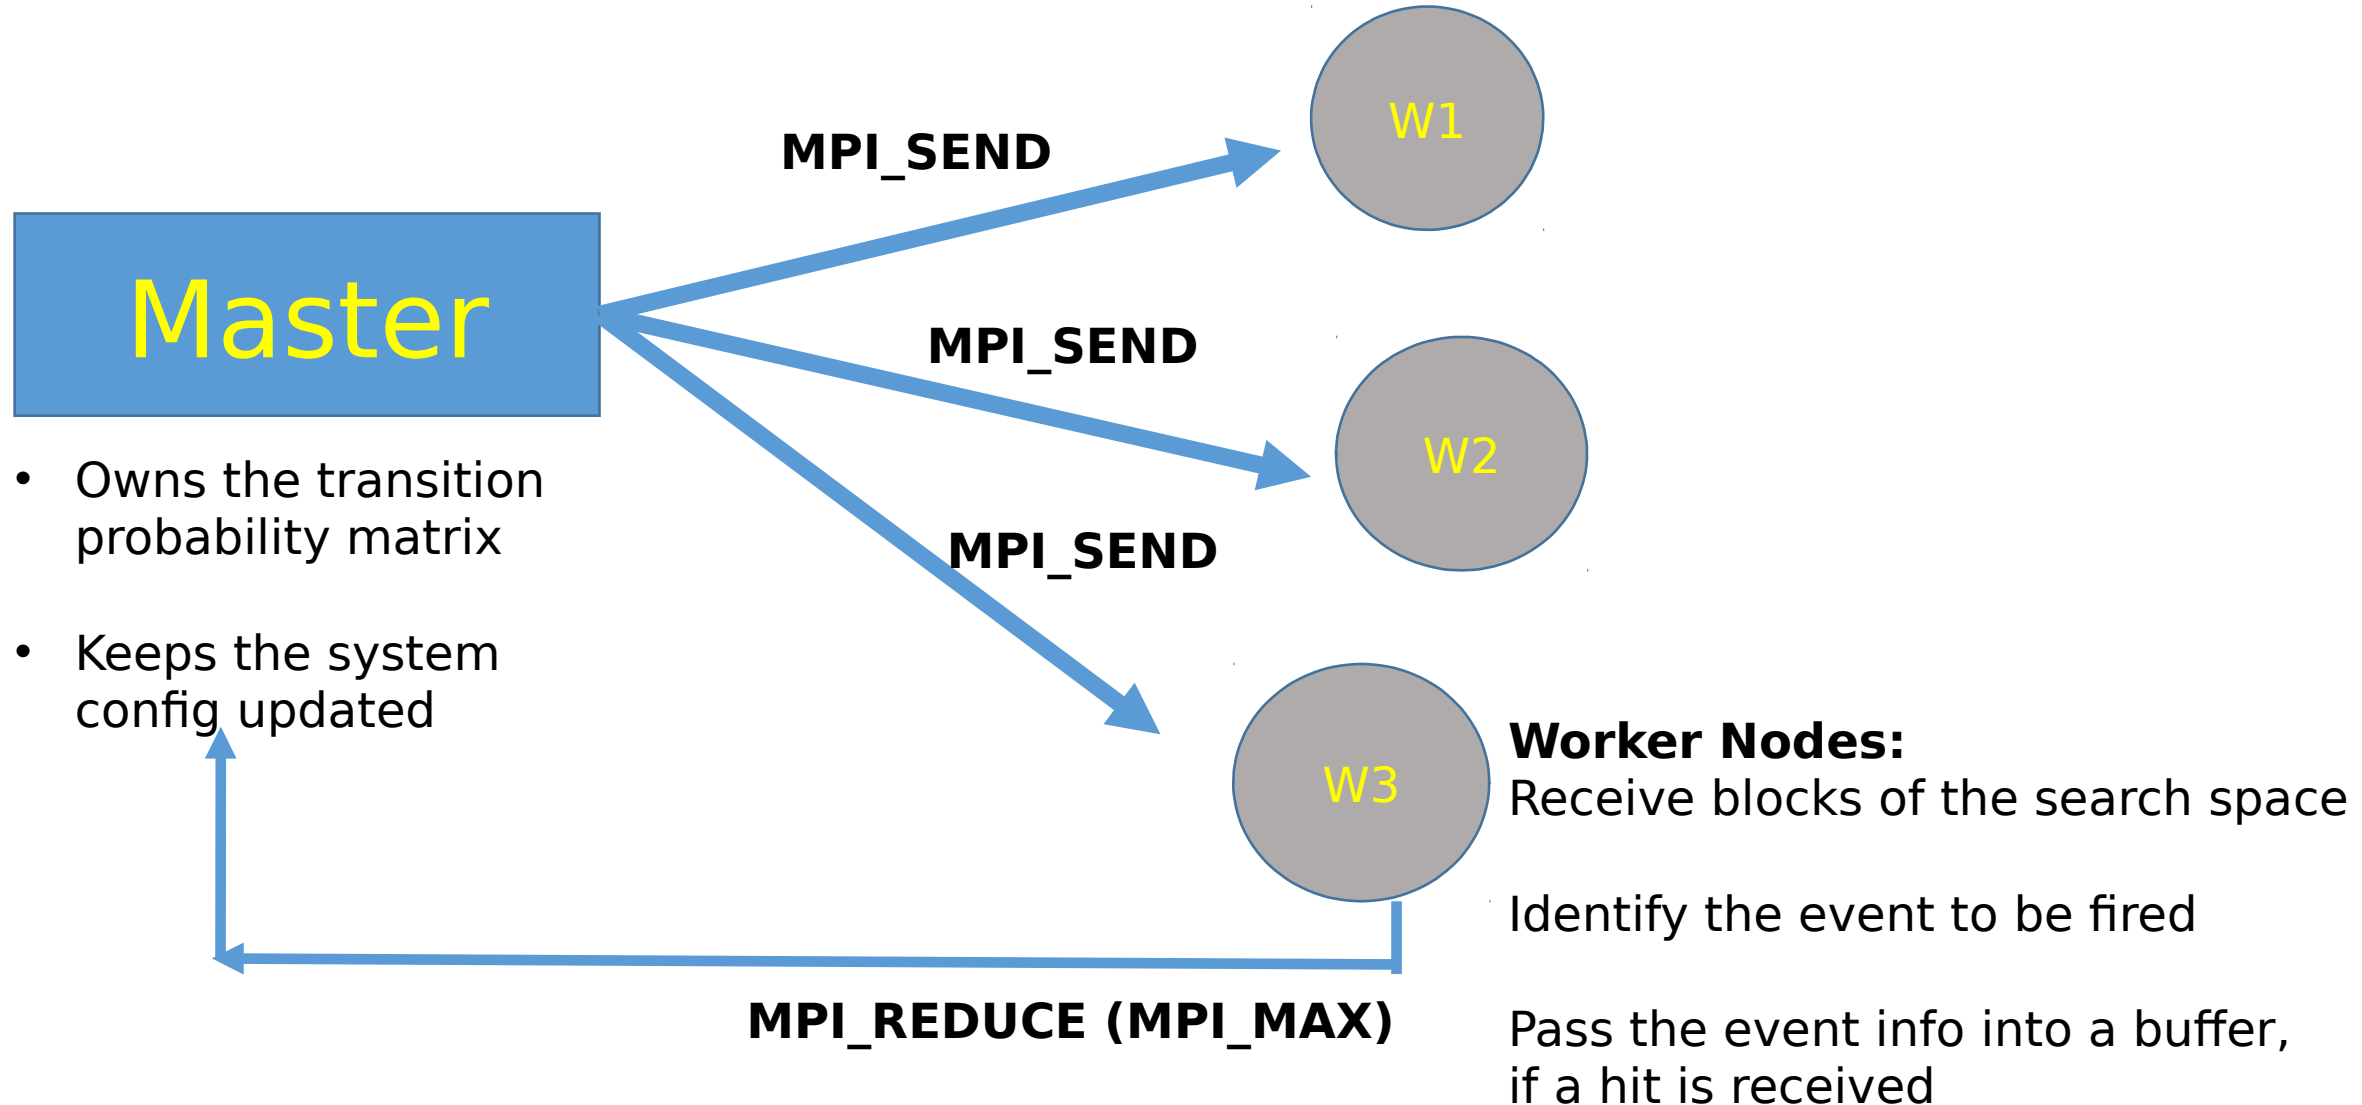




```
179 for(n=0;n<fibnum;n++)
180 {
181     if(n==0)
182     {
183         if ((ranevent < P_sum[4]) && (ranevent > P_init))
184             {
185             if ((P_init < ranevent) && (ranevent < P_sum[n*num_events+0]))
186                 {
187                 fibsize[n] = fibsize[n] + 1;
188                 N_A = N_A - 1;
189                 rccount[n] = rccount[n] + 1;
190                 break;
191                 }
192
193             else if ((P_sum[n*num_events+1] > ranevent) && (ranevent > P_sum[n*num_events+0]))
194                 {
195                 fibsize[n] = fibsize[n] - 1;
196                 N_A = N_A + 1;
197                 rccount[n] = rccount[n] - 1;
198                 break;
199                 }
200
201             else if ((P_sum[n*num_events+2] > (ranevent)) && (ranevent > P_sum[n*num_events+1]))
202                 {
203                 fibsize[n] = fibsize[n] - 1;
204                 betacount[n] = betacount[n] - 1;
205                 N_A = N_A + 1;
206                 break;
207                 }
208
209             else if ((P_sum[n*num_events+3] > (ranevent)) && (ranevent > P_sum[n*num_events+2]))
210                 {
211                 betacount[n] = betacount[n] + 1;
212                 rccount[n] = rccount[n] - 1;
213                 break;

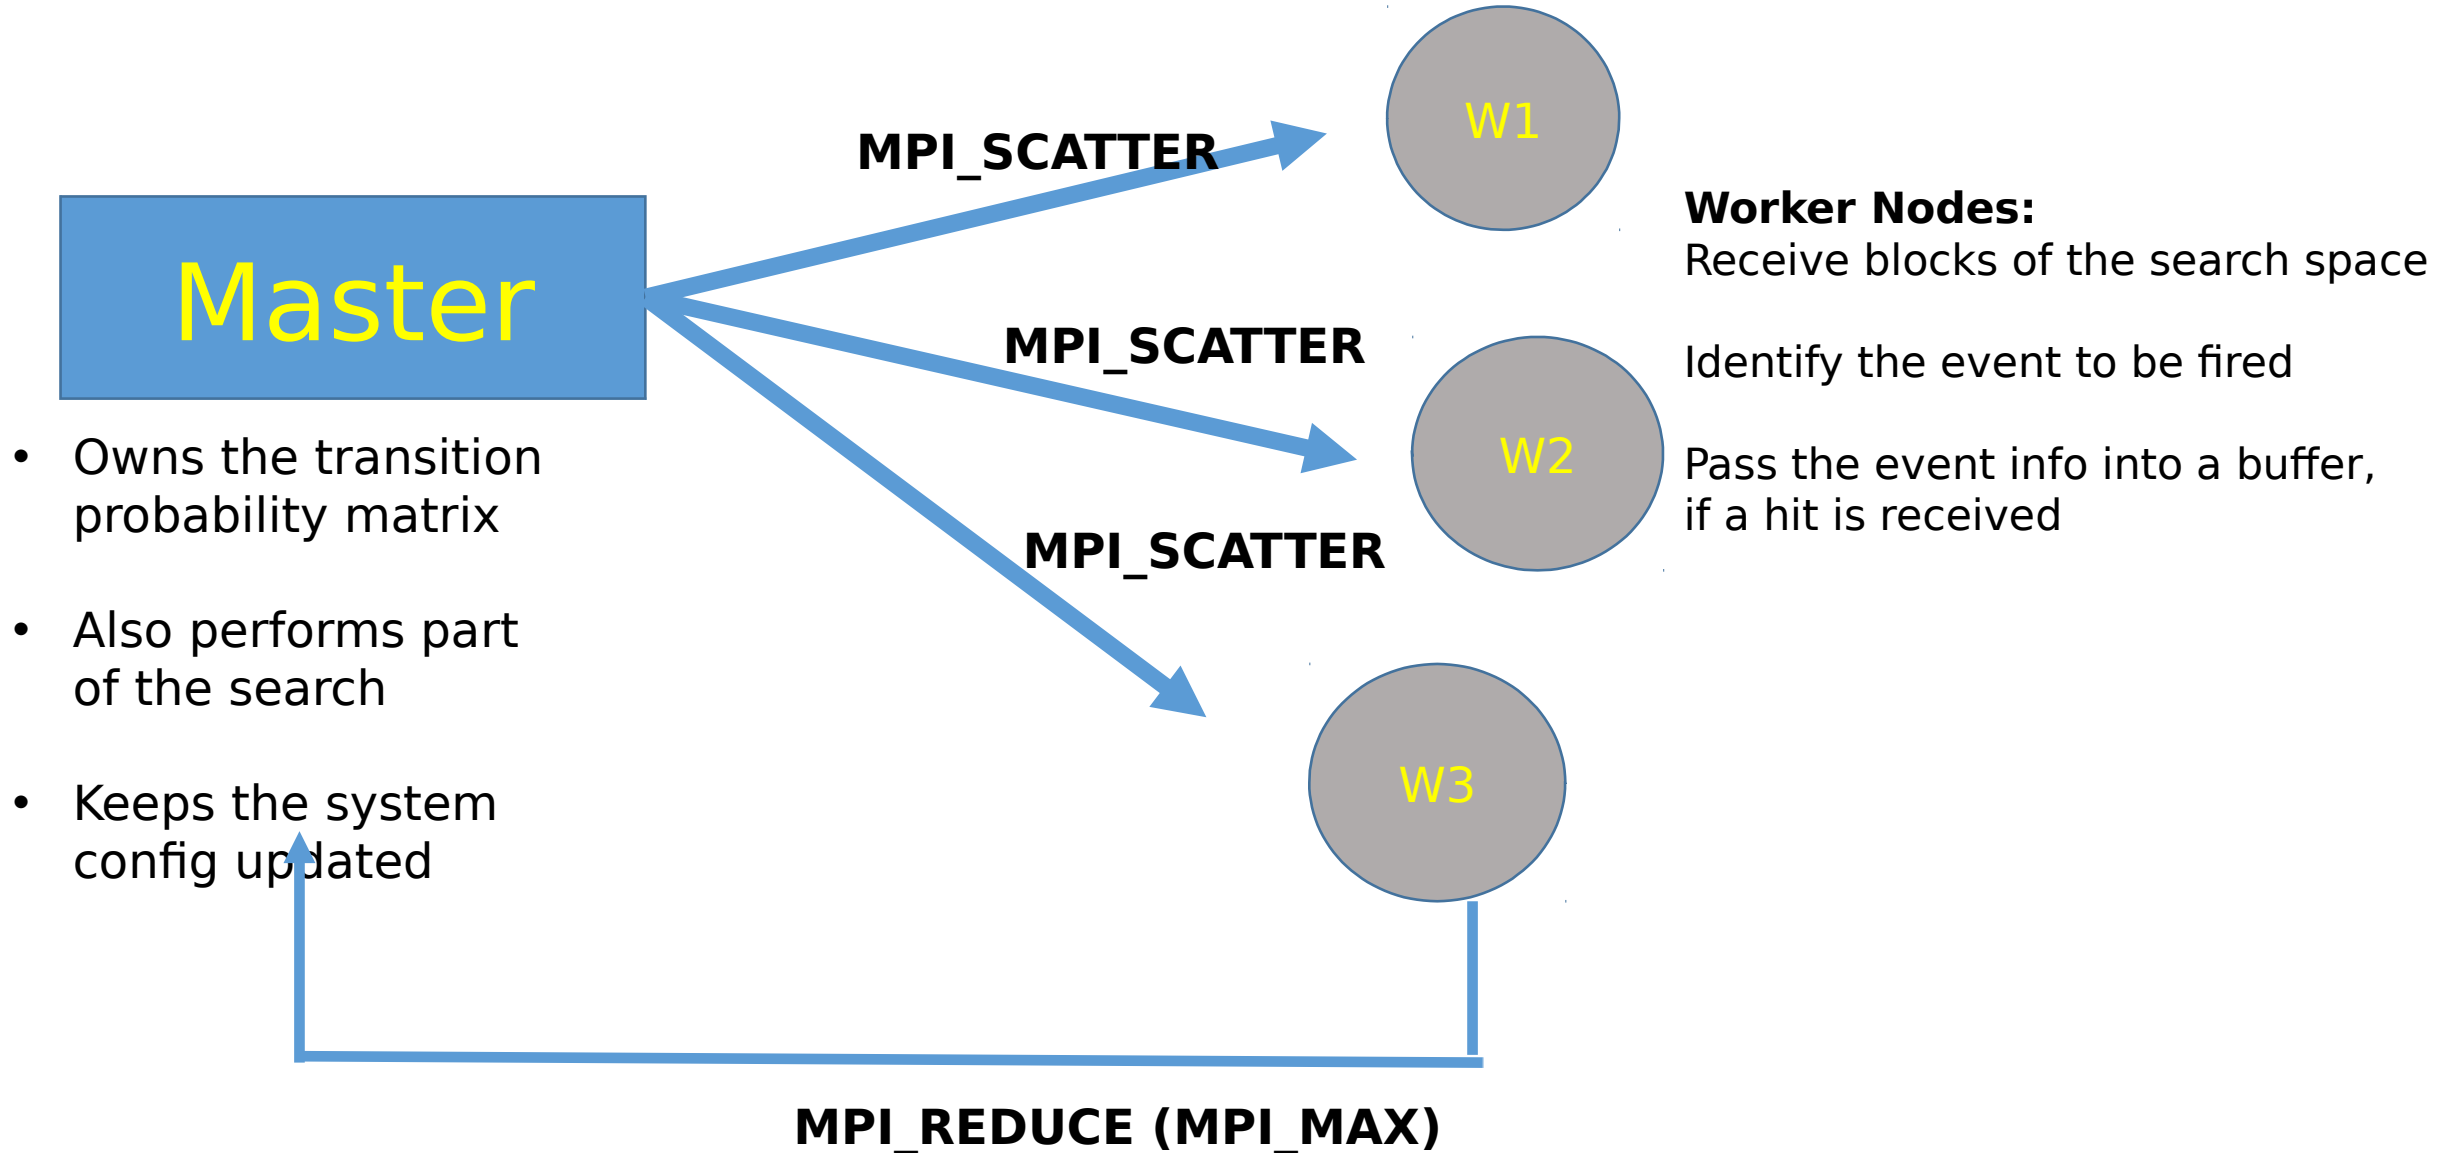
```

# Our attempt (Scheme 1, One-One communications)



```
237     }
238     }
239
240 MPI_Isend(P_send ,num_events*(fibnum/(comm_size-1))+1, MPI_DOUBLE, n_prc, n_prc, MPI_COMM_WORLD,&request);
241
242 n_prc++;
243 }
244 }
245 ranevent = ran2(&id);
246 work_send[0]=0;
247 work_send[1]=0;
248
249     if(my_rank != 0)
250     {
251         MPI_Recv(P_rec ,num_events*(fibnum/(comm_size-1))+1 , MPI_DOUBLE, 0, my_rank, MPI_COMM_WORLD,MPI_STATUS_IGNORE);
252
253         if((ranevent>P_rec[0])&&(ranevent<P_rec[fibnum*num_events/(comm_size-1)]))
254         {
255             for(k=0;k<(fibnum/(comm_size-1));k++)
256             {
257                 for(j=0;j<num_events;j++)
258                 {
259                     if ((P_rec[k*num_events+j] < ranevent) && (ranevent < P_rec[k*num_events+j+1]))
260                     {
261                         event_id=k*num_events+j;
262                         work_send[0]=event_id+1;
263                         work_send[1]=my_rank*((fibnum/(comm_size-1)))-((fibnum/(comm_size-1))-k);
264                         break;
265                     }
266                 }
267             }
268         }
269     }
270 MPI_Barrier( MPI_COMM_WORLD);
271 MPI_Reduce(work_send,work_rec,2,MPI_INT,MPI_MAX,0,MPI_COMM_WORLD);
272
```

# Our attempt (Scheme 2, Collective communication)



```
277 MPI_Scatter(P_sum,num_events*(fibnum/(comm_size)),MPI_DOUBLE,P_rec,num_events*(fibnum/(comm_size)),MPI_DOUBLE,0,MPI_COMM_WORLD);
278 MPI_Scatter>Last_sum,1,MPI_DOUBLE,&Last_rec,1,MPI_DOUBLE,0,MPI_COMM_WORLD);
279
280 ranevent = ran2(&id);
281 work_send[0]=0;
282 work_send[1]=0;
283
284     if(ranevent>Last_rec && ranevent<P_rec[((fibnum/comm_size)*num_events)-1])
285     {
286         for(k=0;k<(fibnum/(comm_size));k++)
287         {
288             for(j=0;(j<num_events);j++)
289             {
290
291                 if(k*num_events+j == 0)
292                 {
293                     if((ranevent>Last_rec) && (ranevent < P_rec[0]))
294                     {
295                         work_send[0]=0;
296                         work_send[1]=my_rank*(fibnum/comm_size)+0;
297                     }
298                 }
299
300                 else if((P_rec[k*num_events+j-1] < ranevent) && (ranevent < P_rec[k*num_events+j]))
301                 {
302                     event_id=k*num_events+j;
303                     work_send[0]=event_id;
304                     work_send[1]=my_rank*(fibnum/comm_size)+k;
305                     break;
306                 }
307             }
308         }
309     }
310
311 MPI_Barrier(MPI_COMM_WORLD);
312 MPI_Reduce(work_send,work_rec,2,MPI_INT,MPI_MAX,0,MPI_COMM_WORLD);
```

# What worked

- Our naïve serial code was optimized to minimize cache misses (a speedup of 1.5 times)
- The MPI code did give us correct results (compared with the serial code and analytical results!)
- Exposed us to a new way of thinking

# What did`nt?

- MPI codes show a speedup from 1 to 3 cores but scale poorly
- Performance slows down at 5 processes or more
- Probably due to huge communication overhead in our code
- Possibly revisit the whole algorithm or use a more parallel-friendly algorithm!

**Thank You!**