


# Parallel Computing Toolkit, Version 2



*Dr. Roman E. Mäder  
MathConsult Dr. R. Mäder  
Samstagernstraße 58a  
CH-8832 Wollerau*

mailto:maeder@mathconsult.ch  
<http://www.mathconsult.ch>

© 2004 MathConsult Dr. R. Mäder. All rights reserved.

- Parallel computation with *Mathematica*
- Experiences with Version 1
- A powerful primitive for parallelization
- Nondeterminism
- Various new features

## Introduction

### | What is Parallel Computing Toolkit?

*Parallel Computing Toolkit* (PCT) is a *Mathematica* AddOn that allows you to develop and run parallel computations.

- distributed memory, master/slave parallelism
- machine independent, all written in *Mathematica*
- uses *MathLink* to communicate with remote kernels
- works on heterogeneous networks, multi-processor machines, LAN and WAN
- scheduling of virtual processes, or explicit distribution to available processors
- virtual shared memory, synchronization, locking
- failure recovery, stranded processes are automatically reassigned
- supports all *Mathematica* data types, including symbolic expressions

## | Parallel Computing ( a System Programmer's View)

Parallel programming requires support for:

- starting processes, waiting for processes to finish.
- scheduling processes on available processors.
- exchanging data between processes, synchronizing access to common resources

In this context:


- a *processor* is a running *Mathematica* kernel. Parallel algorithms are expressed directly in *Mathematica*'s programming language, no compilation is necessary.
- a *process* is an expression to be evaluated. Such an expression may be a simple calculation or a larger program.

## | Requirements

All scheduling happens in the master kernel. No extra software is required on the slave kernels. No support system for parallel computing, such as PVM or MPI is required.

Slave kernels are normally started once and are then used for any number of computations. The Parallel Computing Toolkit contains its own scheduler and load balancing tools.

Master and slave kernels can run on a multi-processor machine, a tightly coupled network of processors (a cluster), a network of workstations or any remote machines accessible through TCP/IP.



In this demonstration the frontend runs on my laptop computer, the master kernel runs on the main node of an Orion Multisystems DT-12, and the slave kernels on the remaining 11 nodes of this desktop cluster.

## | PCT 2

Major new features of Version 2:

- `ParallelEvaluate[ ]` has been extended to become a flexible tool for implementing many structural operations in parallel (such as `Map`, `Apply`, `Cases`, `Select`, `Count`, `MemberQ`, `FreeQ`, `Inner`, `Outer`, and all associative functions).
- `ParallelEvaluate[ ]`, `ParallelMap[ ]`, and `ParallelTable[ ]` now do a single dispatch on each remote kernel, taking relative processor speeds into account for optimal load balancing.
- Delayed definitions for shared variables are now possible, and cause the right sides to be evaluated on the remote kernels.
- Support for different process queue models has been added, including user-defined ones. Standard queues provided are a FIFO queue (as in Version 1) and a new priority queue, as well as a faster unordered queue.
- If no remote kernels are available, all evaluations happen sequentially in the master kernel.

### Enhancements:

- Improved defaults for launching local and remote kernels with a new `Parallel`Configuration`` package.
- `Receive[]` can take a *list* of kernels as argument, and waits for one result from each of them, allowing for callbacks, such as shared variables.
- Aborting remote kernels is now possible, provided the MathLink device used for the kernel connection supports aborts. Resetting runaway kernels is now possible in a wider set of circumstances, in many cases without an actual abort.
- Newly launched kernels inherit all previously exported environments, loaded packages, and declared shared variables.

## | System Configuration

Local configuration, such as a list of commonly available machines and the operating-system commands needed to launch a kernel there, can be set up in a user's `init.m` file.

The template `$RemoteCommand` may contain four slots filled as follows

```
'1'  hostname      first argument of LaunchSlave[]
'2'  linkname      as returned by LinkCreate[]
'3'  remote user   $RemoteUserName, default is $UserName
'4'  protocol      link protocol argument
```

```
(orion) In[1]:= $RemoteCommand
```

```
(orion) Out[1]= ssh -f '1' '/usr/local/bin/math -mathlink -
                linkmode Connect -linkname '2' -noinit >&/dev/null'
```

This is the default list of available remote machines.

```
(orion) In[2]:= $AvailableMachines
```

```
(orion) Out[2]= {n1, n2, n3, n4, n5, n6, n7, n8, n9, n10, n11}
```

## | Startup

```
(orion) In[3]:= Needs["Parallel`Debug`"]
                Needs["Parallel`"]
```

```
Parallel Computing Toolkit 2.0beta3 (October 2004)
```

```
Created by Roman E. Maeder
```

```
(orion) In[6]:= LaunchSlaves[]
```

```
(orion) Out[6]= {slave_1[n1], slave_2[n2], slave_3[n3], slave_4[n4], slave_5[n5], slave_6[n6],
                slave_7[n7], slave_8[n8], slave_9[n9], slave_10[n10], slave_11[n11]}
```

```
(orion) In[7]:= TableForm[RemoteEvaluate[{$ProcessorID, $MachineName, $SystemID,
    $Version, TimeUsed[], MaxMemoryUsed[]}], TableHeadings ->
    {None, {"ID", "host", "OS", "Mathematica Version", "CPU Time", "Memory"}}]
```

```
(orion) Out[7]//TableForm=
  ID      host      OS      Mathematica Version      CPU Time
  1       n1       Linux   5.0 for Linux (November 18, 2003)  0.064991
  2       n2       Linux   5.0 for Linux (November 18, 2003)  0.06499
  3       n3       Linux   5.0 for Linux (November 18, 2003)  0.06599
  4       n4       Linux   5.0 for Linux (November 18, 2003)  0.062989
  5       n5       Linux   5.0 for Linux (November 18, 2003)  0.06599
  6       n6       Linux   5.0 for Linux (November 18, 2003)  0.06899
  7       n7       Linux   5.0 for Linux (November 18, 2003)  0.06799
  8       n8       Linux   5.0 for Linux (November 18, 2003)  0.064991
  9       n9       Linux   5.0 for Linux (November 18, 2003)  0.06199
  10      n10      Linux   5.0 for Linux (November 18, 2003)  0.068989
  11      n11      Linux   5.0 for Linux (November 18, 2003)  0.068989
```

## Parallelizing Commands

### I Wait and Queue

Queue[expr] queues the remote evaluation of expr and returns a process ID (pid).

Wait[pid] waits for the given process to finish and returns its result.

Therefore, Wait[Queue[expr]] is just the evaluation function (or identity).

```
(orion) In[8]:= Wait[Queue[1 + 2]]
```

```
(orion) Out[8]= 3
```

Wait[{pid1, pid2, ...}] waits for several processes.

```
(orion) In[9]:= Wait[Queue /@ {1 + 1, 2 + 2, 3 + 3}]
```

```
(orion) Out[9]= {2, 4, 6}
```

This example was not very parallel. We need to be careful.

```
(orion) In[10]:= Wait[Queue /@ Unevaluated[{1 + 1, 2 + 2, 3 + 3}]]
```

```
(orion) Out[10]= {2, 4, 6}
```

To convince the skeptic.

```
(orion) In[11]:= Wait[Queue /@ Unevaluated[{$ProcessorID, $ProcessorID}]]
```

```
(orion) Out[11]= {1, 2}
```

```
(orion) In[12]:= Wait[Queue /@ {$ProcessorID, $ProcessorID}]
```

```
(orion) Out[12]= {0, 0}
```

### I Elegant, but Slow

A recipe for parallelizing functional operations.

```

(orion) In[13]:= Map[f, {a, b, c}]
(orion) Out[13]= {f[a], f[b], f[c]}

(orion) In[14]:= Composition[q, f][x]
(orion) Out[14]= q[f[x]]

(orion) In[15]:= Map[Composition[q, f], {a, b, c}]
(orion) Out[15]= {q[f[a]], q[f[b]], q[f[c]]}

(orion) In[16]:= Map[Composition[Queue, f], {a, b, c}]
(orion) Out[16]= {f[a], f[b], f[c]}

(orion) In[16]:= Wait[%]

(orion) In[17]:= Wait[Table[Queue[{i}, Prime[i]], {i, 1, 10}]]
(orion) Out[17]= {2, 3, 5, 7, 11, 13, 17, 19, 23, 29}

```

## | A Marketing Problem

In V1, everybody seemed to try out something like this at first.

```

(orion) In[18]:= ParallelTable[Sin[i], {i, 0.0, 100000.0}] // Short
(orion) Out[18]//Short=
{0., 0.841471, <<99998>>, 0.0357488}

```

What is the underlying primitive of parallel evaluation of many *Mathematica* commands?

If we know how many remote processors we have, we do not have to send out 100,000 processes of the form `Sin[i]`, but *one* process of the form `Table[Sin[i], {i, i0, i1}]` to each processor.

## | Types of Parallel Evaluations

```

(orion) In[19]:= {1 + 1, 2 + 2, 3 + 3, 4 + 4}
(orion) Out[19]= {2, 4, 6, 8}

(orion) In[20]:= Map[Prime, {1, 2, 3, 4, 5}]
(orion) Out[20]= {2, 3, 5, 7, 11}

(orion) In[21]:= Select[{1, 2, 3, 4, 5}, EvenQ]
(orion) Out[21]= {2, 4}

(orion) In[22]:= Count[{1, 2, 3, 4, 5}, _?OddQ]
(orion) Out[22]= 3

(orion) In[23]:= {a, b, c} . {x, y, z}
(orion) Out[23]= a x + b y + c z

(orion) In[24]:= Plus[1, 2, 3, 4, 5]
(orion) Out[24]= 15

```

Common theme: iteration over a list (or other long expression).

Efficient parallelization: do not create one process for each element of the list, but work in chunks.

## The New ParallelEvaluate

### | The General Case

The general mechanism to parallelize all such cases:

- Split the input expression into chunks
- Apply a function to each chunk
- Combine the results

Here is ParallelEvaluate:

```
(orion) In[25]:= ?ParallelEvaluate
```

```
ParallelEvaluate[h[exprs...], f, comb] distributes parts of the computation of f[h[exprs]]
to all slaves and combines the partial results with comb. The default combinator
is h if h is Flat, and Join otherwise. The default function f is the Identity.
```

```
(orion) In[26]:= ParallelEvaluate[h[1 + 2, 2 + 3, 3 + 4, 4 + 5, 5 + 6], f, comb]
```

```
(orion) Out[26]= comb[f[h[3]], f[h[5]], f[h[7]], f[h[9]], f[h[11]]]
```

### | Special Case 1: Join

For structure-preserving functions  $f$  the natural combination function is Join.

```
(orion) In[30]:= ParallelEvaluate[{1, 2, 3, 4, 5}, Prime]
```

```
(orion) Out[30]= {2, 3, 5, 7, 11}
```

```
(orion) In[31]:= ParallelEvaluate[{1, 2, 3, 4, 5}, Function[list, Map[f, list]]]
```

```
(orion) Out[31]= {f[1], f[2], f[3], f[4], f[5]}
```

### | Special Case 2: Identity

The default function to perform is the identity (do nothing).

```
(orion) In[32]:= ParallelEvaluate[{1 + 2, 2 + 3, 3 + 4, 4 + 5, 5 + 6}]
```

```
(orion) Out[32]= {3, 5, 7, 9, 11}
```

### | Special Case 3: Associative Functions

For associative  $h$ , the natural parallelization of

$$h[a, b, c, d, e]$$

is

$$h[h[a, b, c], h[d, e]]$$

That is, the function is the Identity, but the combiner is  $h$  itself.

```
(orion) In[33]:= ParallelEvaluate[1 + 2 + 3 + 4 + 5 + 6 + 7 + 8]
(orion) Out[33]= 36
```

## | How to Distribute a Table?

The structure over which to parallelize `Table[expr, iterator]` is the iterator!

```
(orion) In[34]:= Table[i, {i, 1, 7}]
(orion) Out[34]= {1, 2, 3, 4, 5, 6, 7}
```

Of course, we should not generate it explicitly, but figure out suitable subiterators.

```
(orion) In[35]:= Join[Table[i, {i, 1, 4}], Table[i, {i, 5, 7}]]
(orion) Out[35]= {1, 2, 3, 4, 5, 6, 7}
```

Now, parallelization is easy.

```
(orion) In[36]:= ParallelDispatch[Join[Table[i, {i, 1, 4}], Table[i, {i, 5, 7}]]]
(orion) Out[36]= {1, 2, 3, 4, 5, 6, 7}
```

## | Load Balancing

If we know the relative performance of the remote processors, we can take this information into account when figuring out the chunks. Version 2 will allow you to measure and specify this information. This example uses remote kernels of different CPU speeds.

```
TableForm[{ProcessorID[#], #, ProcessorSpeed[#]} & /@$Slaves,
  TableHeadings -> {None, {"ID", "kernel", "speed setting"}}]

ID      kernel                speed setting
1       slave1 [localhost]      1
2       slave2 [sirius]        1
3       slave3 [sirius]        1
4       slave4 [vvv]           1

timings = First /@ RemoteEvaluate[Timing[N[Cos[1], 10000];]] /. Second -> 1
{2.83, 1.59, 1.59, 1.46}

speeds = 1 / timings / Min[1 / timings]
{1., 1.77987, 1.77987, 1.93836}

Round[200 %]
{200, 356, 356, 388}

MapIndexed[(ProcessorSpeed[Extract[$Slaves, #2]] = #1) &, speeds]
{1., 1.77987, 1.77987, 1.93836}
```

## Parallel Commands by the Dozen

### | General Idea

Examine each *Mathematica* command that operates on a list or other long expression.

```
f[list, arg2, arg3] →
ParallelEvaluate[list, Function[li, f[li, arg2, arg3]]]
```

## | Examples

```
(orion) In[37]:= ParallelEvaluate[{1 + 1, 2 + 1, 3 + 1, 4 + 1, 5 + 1},
Function[li, Select[li, EvenQ]]]
```

```
(orion) Out[37]= {2, 4, 6}
```

```
(orion) In[38]:= ParallelEvaluate[{1 + 1, 2 + 1, 3 + 1, 4 + 1, 5 + 1},
Function[li, Cases[li, _?OddQ]]]
```

```
(orion) Out[38]= {3, 5}
```

```
(orion) In[39]:= With[{n = 5},
mata = Table[Random[Real, {-1, 1}], {n}, {n}];
matb = Table[Random[Real, {-1, 1}], {n}, {n}];
]
```

```
(orion) In[40]:= With[{mata = mata, matb = matb},
par = ParallelEvaluate[mata, Dot[#, matb] &]
];
MatrixForm[par]
```

```
(orion) Out[41]//MatrixForm=

$$\begin{pmatrix} -0.413189 & 1.97472 & -0.764574 & 0.398395 & 0.390085 \\ 0.340144 & 1.05595 & 0.140281 & 1.83771 & 0.39134 \\ -0.020901 & 1.99898 & -0.127458 & 0.0551463 & -0.427745 \\ 2.14046 & 0.732813 & 1.41867 & 0.783665 & 0.345912 \\ 0.0672538 & 0.127011 & 0.799246 & -0.172638 & -1.12883 \end{pmatrix}$$

```

## | Less Obvious

For functions that do not preserve the structure, a custom combiner needs to be chosen.

```
(orion) In[42]:= ParallelEvaluate[{1 + 1, 2 + 1, 3 + 1, 4 + 1, 5 + 1}, Count[#, _?OddQ] &, Plus]
```

```
(orion) Out[42]= 2
```

```
(orion) In[43]:= ParallelEvaluate[{1, 2, 3, 4, 5, 6, 7}, MemberQ[#, 5] &, Or]
```

```
(orion) Out[43]= True
```

```
(orion) In[44]:= ParallelEvaluate[{1, 2, 3, 4, 5, 6, 7}, MemberQ[#, 9] &, Or]
```

```
(orion) Out[44]= False
```

```
(orion) In[45]:= ParallelEvaluate[{1, 2, 3, 4, 5, 6, 7}, FreeQ[#, 5] &, And]
```

```
(orion) Out[45]= False
```

```
(orion) In[46]:= ParallelEvaluate[{1, 2, 3, 4, 5, 6, 7}, FreeQ[#, 9] &, And]
```

```
(orion) Out[46]= True
```

## | Associative Functions

```
(orion) In[47]:= ParallelEvaluate[1 + 2 + 3 + 4 + 5 + 6 + 7 + 8]
```

```
(orion) Out[47]= 36
```



```
(orion) In[48]:= ParallelEvaluate[Max[3, 5, 7, 3]]
(orion) Out[48]= 7

(orion) In[49]:= ParallelEvaluate[GCD[2, 4, 6, 8]]
(orion) Out[49]= 2

(orion) In[50]:= ParallelEvaluate[StringJoin["parallel ", "computation ", "is ", "fun."]]
(orion) Out[50]= parallel computation is fun.
```

## | It's Always the Same!

As we saw, to parallelize

```
(orion) In[51]:= Count[{1, 2, 3, 4, 5, 6, 7}, _?PrimeQ]
(orion) Out[51]= 4
```

you have to say

```
(orion) In[52]:= ParallelEvaluate[{1, 2, 3, 4, 5, 6, 7}, Count[#, _?PrimeQ] &, Plus]
(orion) Out[52]= 4
```

The idea is always the same. Wouldn't this be nice:

```
(orion) In[53]:= ParallelEvaluate[Count[{1, 2, 3, 4, 5, 6, 7}, _?PrimeQ]]
(orion) Out[53]= 4
```

It is quite simple:

```
Count /: ParallelEvaluate[Count[list_, params___]] :=
  ParallelEvaluate[list, Count[#, params] &, Plus]
```

All of these have been overloaded.

```
(orion) In[54]:= $ParallelCommands
(orion) Out[54]= {Cases, Select, Count, MemberQ, FreeQ,
  Map, Apply, Outer, Inner, Dot, Table, Sum, Product}
```

## | The Ultimate Goal

It is still a long way from

```
ParallelEvaluate[your program goes here ...]
```

# Nondeterminism

## | Waiting for a Result

Idea: try out different methods to solve a problem. The first one to succeed wins, the others are then aborted. We can also allow for one method to return `$Failed` to indicate that it did not succeed. Waiting then continues until one result is found that is other than `$Failed`.

It is already supported in V1. Queue each evaluation as a process, then use `WaitOne[ ]`. This returns with one of the results. You can then abort the others (this part does not work well in V1).

```
(orion) In[55]:= ids = {Queue[1 + 1], Queue[2 + 2]}
(orion) Out[55]= {pid25, pid26}
(orion) In[56]:= {res, id, ids} = WaitOne[ids]
(orion) Out[56]= {2, pid25, {pid26}}
(orion) In[57]:= ResetSlaves[];
(orion) In[58]:= res
(orion) Out[58]= 2
```

## | Parallel Try

Here is a function that combines these steps.

The function `ParallelTry[{exprs...}]` takes a list of expressions as argument and submits them to remote kernels. The first result other than `$Failed` is returned, and any remaining computations are aborted.

```
ParallelTry[exprs_] :=
Module[{res = $Failed, id, ids},
  ids = List @@ Queue /@ Unevaluated[exprs];
  While[res === $Failed && Length[ids] > 0, (* try one more *)
    {res, id, ids} = WaitOne[ids]
  ];
  ResetSlaves[];
  res
]
```

## | Finding Irreducible Polynomials

We want to find an irreducible polynomial modulo  $n$ .

This generates a random polynomial of a given degree mod  $p$ .

```
p = 2 (*the modulus *)
2
randPoly[n_Integer?Positive, x_, p_] :=
x^n + Sum[Random[Integer, {0, p - 1}] * x^i, {i, 1, n - 1}] +
  Random[Integer, {1, p - 1}]
```

It is irreducible if it does not have any factors other than a numerical content.

```
irreducibleQ[poly_, p_] :=
  With[{fl = FactorList[poly, Modulus -> p]}, Length[fl] < 3 && fl[[-1, 2]] == 1]
poly = randPoly[10, x, p]
1 + x3 + x8 + x10
irreducibleQ[poly, p]
False
```

## | The Loop

This loops keeps trying until it finds an irreducible one.

```

While[! irreducibleQ[poly = randPoly[100, x, p], p], Null]; poly
1 + x + x2 + x7 + x8 + x10 + x12 + x14 + x15 + x16 + x17 + x19 + x20 + x22 + x24 + x26 +
x28 + x29 + x33 + x34 + x35 + x36 + x37 + x39 + x43 + x44 + x45 + x46 + x47 + x48 +
x49 + x50 + x51 + x52 + x56 + x57 + x58 + x59 + x61 + x62 + x63 + x68 + x69 + x70 +
x72 + x74 + x76 + x78 + x80 + x82 + x86 + x87 + x90 + x91 + x92 + x94 + x95 + x97 + x100

```

Here is the function that combines these steps.

```

findOne[n_, p_] :=
Module[{poly},
  While[! irreducibleQ[poly = randPoly[n, x, p], p], Null];
  poly]

findOne[100, p]

1 + x2 + x4 + x5 + x7 + x9 + x10 + x11 + x12 + x15 + x17 + x18 + x21 + x22 +
x25 + x26 + x27 + x33 + x34 + x36 + x37 + x38 + x39 + x43 + x45 + x46 + x47 +
x48 + x54 + x56 + x60 + x61 + x62 + x64 + x66 + x71 + x72 + x75 + x76 + x77 +
x78 + x81 + x83 + x84 + x91 + x92 + x93 + x94 + x95 + x96 + x97 + x98 + x100

```

## | In Parallel

```
ExportEnvironment[randPoly, irreducibleQ, findOne]
```

Generate list of identical commands, one for each remote kernel.

```

With[{n = 512, p = 2},
  Join @@ Table[Hold[findOne[n, p]], {Length[$Slaves]}]
]

Hold[findOne[1000, 2], findOne[1000, 2],
  findOne[1000, 2], findOne[1000, 2], findOne[1000, 2]]

```

Let them loose.

## ParallelTry[%]

```

1 + x2 + x3 + x5 + x11 + x14 + x15 + x17 + x20 + x22 + x23 + x25 + x28 + x29 + x34 + x39 +
x40 + x41 + x43 + x44 + x45 + x47 + x49 + x50 + x53 + x54 + x56 + x57 + x58 + x59 + x62 +
x64 + x65 + x68 + x69 + x83 + x86 + x87 + x98 + x101 + x102 + x105 + x109 + x110 + x111 +
x112 + x113 + x114 + x115 + x116 + x118 + x119 + x121 + x125 + x126 + x127 + x128 + x130 +
x131 + x134 + x136 + x139 + x141 + x143 + x145 + x147 + x150 + x151 + x154 + x155 + x158 +
x162 + x164 + x165 + x167 + x168 + x170 + x171 + x173 + x177 + x179 + x180 + x181 + x182 +
x184 + x185 + x190 + x192 + x195 + x197 + x198 + x200 + x201 + x203 + x204 + x205 + x207 +
x208 + x211 + x213 + x214 + x215 + x219 + x220 + x224 + x226 + x230 + x232 + x233 + x234 +
x235 + x238 + x239 + x240 + x242 + x244 + x245 + x248 + x249 + x250 + x252 + x255 + x256 +
x260 + x261 + x262 + x264 + x265 + x266 + x267 + x269 + x271 + x274 + x275 + x277 + x279 +
x280 + x283 + x284 + x285 + x288 + x290 + x291 + x294 + x297 + x298 + x299 + x300 + x301 +
x302 + x304 + x305 + x306 + x308 + x312 + x314 + x318 + x319 + x322 + x326 + x329 + x330 +
x333 + x335 + x336 + x338 + x341 + x342 + x344 + x345 + x346 + x348 + x349 + x352 + x353 +
x354 + x355 + x357 + x358 + x362 + x364 + x366 + x368 + x371 + x372 + x374 + x377 + x378 +
x379 + x380 + x382 + x383 + x384 + x385 + x387 + x392 + x394 + x397 + x399 + x400 + x403 +
x404 + x406 + x409 + x414 + x415 + x418 + x419 + x420 + x421 + x427 + x439 + x440 + x441 +
x442 + x443 + x447 + x448 + x451 + x452 + x453 + x455 + x456 + x459 + x462 + x464 + x465 +
x466 + x467 + x468 + x473 + x474 + x475 + x479 + x480 + x481 + x482 + x483 + x485 + x490 +
x493 + x494 + x496 + x497 + x498 + x499 + x500 + x501 + x502 + x503 + x505 + x506 + x509 +
x510 + x512 + x514 + x515 + x517 + x518 + x519 + x521 + x522 + x526 + x531 + x532 + x533 +
x534 + x536 + x537 + x538 + x541 + x545 + x546 + x547 + x549 + x550 + x551 + x552 + x554 +
x555 + x556 + x557 + x558 + x561 + x566 + x567 + x570 + x571 + x572 + x574 + x575 + x577 +
x584 + x585 + x586 + x588 + x589 + x593 + x594 + x600 + x603 + x604 + x605 + x607 + x609 +
x612 + x613 + x614 + x615 + x618 + x620 + x623 + x625 + x630 + x632 + x633 + x634 + x638 +
x642 + x644 + x646 + x648 + x649 + x652 + x655 + x656 + x659 + x660 + x661 + x662 + x663 +
x664 + x665 + x666 + x667 + x668 + x669 + x670 + x671 + x672 + x673 + x674 + x679 + x681 +
x682 + x683 + x684 + x685 + x687 + x688 + x689 + x692 + x695 + x697 + x702 + x703 + x704 +
x706 + x707 + x708 + x710 + x714 + x719 + x720 + x724 + x725 + x731 + x732 + x734 + x741 +
x748 + x749 + x751 + x752 + x753 + x754 + x755 + x759 + x766 + x767 + x769 + x771 + x773 +
x775 + x777 + x782 + x783 + x784 + x785 + x790 + x791 + x792 + x793 + x795 + x797 + x798 +
x799 + x800 + x802 + x803 + x804 + x805 + x806 + x808 + x810 + x811 + x812 + x813 + x815 +
x816 + x817 + x818 + x819 + x821 + x822 + x823 + x828 + x829 + x830 + x831 + x833 + x835 +
x836 + x837 + x840 + x842 + x843 + x844 + x846 + x847 + x848 + x850 + x853 + x854 + x855 +
x856 + x858 + x863 + x865 + x866 + x867 + x868 + x869 + x870 + x871 + x872 + x873 + x877 +
x878 + x879 + x884 + x887 + x888 + x890 + x891 + x892 + x893 + x896 + x898 + x900 + x901 +
x902 + x903 + x905 + x907 + x909 + x910 + x911 + x913 + x916 + x917 + x919 + x923 + x925 +
x926 + x928 + x929 + x931 + x932 + x933 + x934 + x935 + x937 + x939 + x940 + x943 + x947 +
x949 + x952 + x957 + x961 + x962 + x963 + x965 + x968 + x969 + x970 + x972 + x973 + x978 +
x979 + x980 + x983 + x984 + x985 + x986 + x988 + x989 + x990 + x992 + x996 + x998 + x1000

```

(In this example, 853 polynomials were examined.)

## I Comparing Algorithms

This example from the new collection of sample applications shows how to compare the different methods for numerical minimization available in *Mathematica* V5.

$$\begin{aligned}
 f[x_, y_] &:= e^{\sin[50x]} + \sin[60e^y] + \\
 &\quad \sin[70 \sin[x]] + \sin[\sin[80y]] - \sin[10(x+y)] + \frac{1}{4}(x^2 + y^2)
 \end{aligned}$$

We can run all methods in parallel to compare the results.

```
TableForm[ParallelMap[tryMethod, methods], TableDepth -> 2,
  TableHeadings -> {None, {Method, Timing, Minimum, Position}}]

```

Method	Timing	Minimum	Position
DifferentialEvolution	0.51	0.998004	{x -> -31.9783, y -> -31.97}
NelderMead	0.46	16.4409	{x -> -15.9634, y -> 15.963}
RandomSearch	1.11	0.998004	{x -> -31.9783, y -> -31.97}
SimulatedAnnealing	1.77	0.998004	{x -> -31.9783, y -> -31.97}
Automatic	0.38	15.5038	{x -> -31.9317, y -> 15.965}

Here we use `ParallelTry` to find the fastest method and abort all other calculations. There is no point in including the `Automatic` method setting, because it implies on of the other methods.

```

trials = tryMethod /@ Hold @@ Complement[methods, {Automatic}];

ParallelTry @@ trials

{NelderMead, 0.38, 16.4409, {x -> -15.9634, y -> 15.9634}}

```

## Shared Memory

### | Communication

How can processes communicate with each other?

- Message Passing
- Shared Memory

*Mathematica* kernels do not share memory, even if they run on the same machine. Our computation model is always a distributed memory one.

**Virtual Shared Memory:** implementing shared memory on distributed memory machines with a software layer that handles all variable accesses, using message passing.

When a process needs the value of a global (shared) variable or wants to set its value, it sends a request back to the scheduler. The scheduler maintains the variables, updates their values, and sends back the requested values.

### | Examples

We will work with two slave kernels.

```
(orion) In[59]:= s1 = $Slaves[[1]]; s2 = $Slaves[[2]];
```

Declare shared variables.

```
(orion) In[60]:= SharedVariables[{x}]
```

Basic tests. One slave kernel reads the variable, another one increments its value.

```
(orion) In[61]:= x = 0;
```

```
(orion) In[62]:= RemoteEvaluate[x = x + 1, s1]
```

```
(orion) Out[62]= 1
```

```
(orion) In[63]:= RemoteEvaluate[x, s2]
```

```
(orion) Out[63]= 1
```

The new value is stored globally.

```
(orion) In[64]:= x
(orion) Out[64]= 1
```

Increment the value in one indivisible step. This operation is important for synchronization.

```
(orion) In[65]:= RemoteEvaluate[x++, s2]
(orion) Out[65]= 1
(orion) In[66]:= RemoteEvaluate[x, s1]
(orion) Out[66]= 2
```

No special syntax is required to use shared variables. Just do things the normal way. If a variable is shared, all accesses are transparently sent to the master process.

```
(orion) In[67]:= ClearShared[x];
```

## | Synchronization

Exclusive access to a variable or other resource can be controlled by a *lock*. An update of a shared variable implements a lock and allows semaphores and critical sections.

Here is the typical code that implements a *critical section* with a lock. Each of the five parallel processes accesses the same shared variable *y*, first reading it, then writing a new value. Before doing this, it acquires a lock, which is released after the critical section. In this way no conflicting updates happen, as can be seen by the output, in which all values from 1 to 5 occur.

```
(orion) In[68]:= SharedVariables[{y, lock}]
(orion) In[69]:= y = 0;
ParallelMap[Module[{a},
  Pause[0.2 Random[]];
  While[TestAndSet[lock, #] != #, Pause[0.1 Random[]]];
  a = y;
  Pause[Random[]];
  y = a + 1;
  lock = Null;
  a + 1] &,
  Range[5]]
(orion) Out[70]= {5, 4, 3, 2, 1}
(orion) In[71]:= ClearSlaves[];
```

## | Shared Downvalues

Delayed definitions (downvalues) of shared variables can be evaluated on the master kernel or the remote kernel that asked for it. Shared downvalues can be used to access shared resources such as a global queue.

## Debugging News

### | Tracing Events

```
(orion) In[72]:= Options[$DebugObject]
(orion) Out[72]= {Trace -> {}, TraceHandler -> Print}
```

```
(orion) In[73]:= OptionValues[Trace]
(orion) Out[73]= {MathLink, Queueing, SendReceive, SharedMemory}

(orion) In[74]:= SetOptions[$DebugObject, Trace → {SendReceive, Queueing, SharedMemory}]
(orion) Out[74]= {Trace → {Queueing, SendReceive, SharedMemory}}

(orion) In[75]:= SetOptions[$DebugObject, TraceHandler → Save]
(orion) Out[75]= {TraceHandler → Save}
```

## | Running A Computation with Tracing

```
(orion) In[76]:= SharedVariables[x]; x = 5;
(orion) In[77]:= Wait[{Queue[x = x + 1], Queue[x = x + 1]}]
(orion) Out[77]= {6, 6}
```

## | Analyzing the Trace

```
(orion) In[78]:= TableForm[TraceList[], TableDepth → 2,
  TableHeadings → {Automatic, {"Trigger", "Event"}}]

(orion) Out[78]//TableForm=
      Trigger          Event
1      SendReceive    Sending to slave1[n1]: -internal value- (q=1)
2      SendReceive    Sending to slave2[n2]: -internal value- (q=1)
3      SendReceive    Sending to slave3[n3]: -internal value- (q=1)
4      SendReceive    Sending to slave4[n4]: -internal value- (q=1)
5      SendReceive    Sending to slave5[n5]: -internal value- (q=1)
6      SendReceive    Sending to slave6[n6]: -internal value- (q=1)
7      SendReceive    Sending to slave7[n7]: -internal value- (q=1)
8      SendReceive    Sending to slave8[n8]: -internal value- (q=1)
9      SendReceive    Sending to slave9[n9]: -internal value- (q=1)
10     SendReceive    Sending to slave10[n10]: -internal value- (q=1)
11     SendReceive    Sending to slave11[n11]: -internal value- (q=1)
12     SendReceive    Receiving from slave1[n1]: -internal value- (q=0)
13     SendReceive    Receiving from slave2[n2]: -internal value- (q=0)
14     SendReceive    Receiving from slave3[n3]: -internal value- (q=0)
15     SendReceive    Receiving from slave4[n4]: -internal value- (q=0)
16     SendReceive    Receiving from slave5[n5]: -internal value- (q=0)
17     SendReceive    Receiving from slave6[n6]: -internal value- (q=0)
18     SendReceive    Receiving from slave7[n7]: -internal value- (q=0)
19     SendReceive    Receiving from slave8[n8]: -internal value- (q=0)
20     SendReceive    Receiving from slave9[n9]: -internal value- (q=0)
21     SendReceive    Receiving from slave10[n10]: -internal value- (q=0)
22     SendReceive    Receiving from slave11[n11]: -internal value- (q=0)
23     Queueing       pid1 queued (0)
24     Queueing       pid2 queued (1)
25     SendReceive    Sending to slave1[n1]: pid1[x = x + 1] (q=1)
26     Queueing       pid1 on slave1[n1]
27     SendReceive    Sending to slave2[n2]: pid2[x = x + 1] (q=1)
28     Queueing       pid2 on slave2[n2]
29     SharedMemory   slave1[n1]: x → 5
30     SharedMemory   slave2[n2]: x → 5
31     SharedMemory   slave1[n1]: x = 6 → 6
32     SharedMemory   slave2[n2]: x = 6 → 6
33     SendReceive    Receiving from slave1[n1]: pid1[6] (q=0)
34     Queueing       pid1 done
35     SendReceive    Receiving from slave2[n2]: pid2[6] (q=0)
36     Queueing       pid2 done
37     Queueing       pid1 dequeued
38     Queueing       pid2 dequeued

(orion) In[79]:= newTraceList[]

(orion) In[80]:= SetOptions[$DebugObject, TraceHandler → Print, Trace → {}]

(orion) Out[80]= {Trace → {}, TraceHandler → Print}

(orion) In[81]:= ClearSlaves[];
```

## The New Examples Directory

From Application Examples, an annotated list of all new examples in the Examples subdirectory:



- Performance Measurement and Calibration presents ideas for measuring performance of a parallel computation and shows how to calibrate remote processor speeds for better scheduling.
- Interval Minimum presents a distributed global minimization using interval arithmetic.
- Diameter of a Point Set gives a parallel solution to a classic programming problem. This example is of a tutorial nature.
- Single-Image Stereograms shows how an existing sequential program can be parallelized with a few small modifications.
- Advanced Use of Shared Variables presents various advanced topics on the use of virtual shared memory.
- Parallel Try shows how to set up nondeterministic computations.
- Large Calculations gives examples of large, infinite, or batch computations.
- Graphics and Animations provides examples for using the `Parallel`Commands`` package.
- The Cascade explains how to evaluate associative expressions using a binary tree.

## Summary

- Parallel computation with *Mathematica*
- Distribution of loop-like *Mathematica* commands
- Automatic Parallelization
- Nondeterministic Parallelism