

# Paralelizace výpočtů v systému *Mathematica*

Zdeněk Buk

---

bukz1@fel.cvut.cz

České vysoké učení technické v Praze  
Fakulta elektrotechnická  
Katedra počítačů

2010



# Úvod

## Obsah prezentace

### Paralelizace v systému *Mathematica* obecně

### Novinky ve verzi 7

- Technologie
- Licencování a terminologie

### Paralelizace v systému *Mathematica* konkrétně

- Technologie LightweightGrid

### Další možnosti, stručně

- *MathLink*, nVidia CUDA



# Úvod Paralelizace v systému *Mathematica* obecně

## Paralelizace na úrovni jazyka *Mathematica*

- Netřeba nízkoúrovňového programování
- Odpadá řešení komunikační vrstvy, sdílení dat apod.
- Snadná úprava stávajícího kódu

```
Parallelize[  
  Map[f, data]  
]
```

## Paralelizace na nižší úrovni

- V případě potřeby je možné využít nízkoúrovňových funkcí pro vlastní distribuci částí výpočtů

```
ParallelSubmit[1+2]  
WaitAll[%]
```

## Jak?

- *Mathematica* 6 a nižší - knihovna **Parallel Computing Toolkit**
- Od verze *Mathematica* 7 odpadá nutnost knihovny, podpora pro paralelní výpočty je součástí základní instalace.



# Mathematica 7 vs. 6

## Mathematica $\leq$ 6

- Front End
- Kernel

## Mathematica 7

- Front End
- Control Kernel
- Compute Kernel (*Komunikuje pouze s řídicím jádrem - Control Kernel*)

## Technologie

- Podpora pro paralelní výpočty přímo v systému *Mathematica*
- Odpadá nutnost knihovny Parallel Computing Toolkit
- Automatická detekce výpočetních jader (*Lightweight Grid*)

# Podpora pro automatickou paralelizaci výpočtů



## **Mathematica 7** Ukázka licencování

- *Mathematica 7* je optimalizovaná pro použití na čtyřjádrových počítačích.

### **Mathematica Single Machine license**

- 2 × *Mathematica* Front End
- 2 × Control Kernel (řídící jádro)
- 4 × Compute Kernel (výpočetní jádro)

### **Mathematica Network Increment**

- 1 × *Mathematica* Front End
- 1 × Control Kernel (řídící jádro)
- 4 × Compute Kernel (výpočetní jádro)



## Mathematica 7 Ukázka použití

### Konfigurace a spuštění výpočetních jader

- Přehled dostupných služeb (jader) a licencí
- `$ConfiguredKernels`, `$MaxLicenseProcesses`, `$MaxLicenseSubprocess`

In[1]:=

**`$ConfiguredKernels`**

```
{LightweightGridClient`LightweightGrid[{Agent →
  http://stroj1.domena:3737/WolframLightweightGrid/
  Manager, KernelCount → 4,
  LocalLinkMode → Connect, Service → , Timeout → 5}],
LightweightGridClient`LightweightGrid[
  {Agent →
  http://stroj2.domena:3737/WolframLightweightGrid/
  Manager, KernelCount → 8,
  LocalLinkMode → Connect, Service → , Timeout → 5}],
<<2 local kernels>> }
```

```
In[4]:= $MaxLicenseProcesses
```

```
Out[4]= 8
```

```
In[5]:= $MaxLicenseSubprocesses
```

```
Out[5]= 16
```

- Spuštění, přehled běžících a ukončení běhu výpočetních jader

```
In[6]:= LaunchKernels [{"localhost", "localhost"}]
```

```
Out[6]= {KernelObject[1, local], KernelObject[2, local]}
```

```
In[7]:= Kernels []
```

```
Out[7]= {KernelObject[1, local], KernelObject[2, local]}
```

```
In[8]:= CloseKernels []
```

```
Out[8]= {KernelObject[1, local, <defunct>],  
KernelObject[2, local, <defunct>]}
```

---

## Jednoduché vyhodnocování v paralelním prostředí

```
In[9]:= LaunchKernels [{"localhost", "localhost"}]
```

```
Out[9]= {KernelObject[3, local], KernelObject[4, local]}
```

```
In[10]:= Kernels []
```

```
Out[10]=
```

```
{KernelObject[3, local], KernelObject[4, local]}
```

```
In[11]:= LaunchKernels [{"localhost", "localhost"}]
```

```
Out[11]=
```

```
{KernelObject[5, local], KernelObject[6, local]}
```

```
In[12]:= Kernels []
```

```
Out[12]=
```

```
{KernelObject[3, local], KernelObject[4, local],  
KernelObject[5, local], KernelObject[6, local]}
```

- Vyhodnocení výrazu v paralelním prostředí (*vyhodnocení stejného výrazu na všech výpočetních jádrech*)

```
In[13]:= ParallelEvaluate [$KernelID]
```

```
Out[13]=
```

```
{3, 4, 5, 6}
```

```
In[14]:= ParallelEvaluate[{$KernelID, 1 + 1}]
```

```
Out[14]=
```

```
{{3, 2}, {4, 2}, {5, 2}, {6, 2}}
```

```
In[15]:= ParallelEvaluate[{$KernelID, RandomInteger[{1, 10 000}]}]
```

```
Out[15]=
```

```
{{3, 8348}, {4, 4325}, {5, 6114}, {6, 7355}}
```

```
In[16]:= ParallelEvaluate [ { $KernelID, $ProcessID, $OperatingSystem,  
    $MachineType, $Version } ]
```

Out[16]=

```
{ { 3, 7016, MacOSX, PC,  
    7.0 for Mac OS X x86 (64-bit) (February 19, 2009) },  
  { 4, 7017, MacOSX, PC,  
    7.0 for Mac OS X x86 (64-bit) (February 19, 2009) },  
  { 5, 7018, MacOSX, PC,  
    7.0 for Mac OS X x86 (64-bit) (February 19, 2009) },  
  { 6, 7019, MacOSX, PC,  
    7.0 for Mac OS X x86 (64-bit) (February 19, 2009) } }
```

```
In[17]:= CloseKernels [ ]
```

Out[17]=

```
{ KernelObject [ 3, local, <defunct> ],  
  KernelObject [ 4, local, <defunct> ],  
  KernelObject [ 5, local, <defunct> ],  
  KernelObject [ 6, local, <defunct> ] }
```

## Automatická paralelizace

```
In[18]:= LaunchKernels [{"localhost", "localhost"}]
```

Out[18]=

```
{KernelObject[7, local], KernelObject[8, local]}
```

```
In[19]:= AbsoluteTiming[Select[Range[4000], PrimeQ[2# - 1] &]]
```

Out[19]=

```
{11.766954, {2, 3, 5, 7, 13, 17, 19, 31, 61,  
89, 107, 127, 521, 607, 1279, 2203, 2281, 3217}}
```

```
In[20]:= AbsoluteTiming[  
  Parallelize[Select[Range[4000], PrimeQ[2# - 1] &]]]
```

Out[20]=

```
{7.741839, {2, 3, 5, 7, 13, 17, 19, 31, 61,  
89, 107, 127, 521, 607, 1279, 2203, 2281, 3217}}
```

```
In[21]:= CloseKernels [ ]
```

```
Out[21]=
```

```
{KernelObject[7, local, <defunct>],  
 KernelObject[8, local, <defunct>]}
```



## Výpočty ParallelEvaluate, ParallelMap, ...

```
In[22]:= LaunchKernels [{"localhost", "localhost"}]
```

```
Out[22]=
```

```
{KernelObject[9, local], KernelObject[10, local]}
```

### Jednoduché vyhodnocení

```
In[23]:= ParallelEvaluate[$KernelID]
```

```
Out[23]=
```

```
{9, 10}
```

### Proměnné

```
In[24]:= x = 2;
```

```
In[25]:= ParallelEvaluate[x === 2]
```

```
Out[25]=
```

```
{False, False}
```

```
In[26]:= ParallelEvaluate[x == 2]
```

```
Out[26]=
```

```
{True, True}
```

```
In[27]:= With[{a = 2}, ParallelEvaluate[a === 2] ]
```

```
Out[27]=
```

```
{True, True}
```

```
In[28]:= kernel = First@Kernels []
```

```
Out[28]=
```

```
KernelObject[9, local]
```

```
In[29]:= Table[ParallelEvaluate[IntegerQ[i], kernel], {i, 1, 10}]
```

```
Out[29]=
```

```
{False, False, False, False,  
 False, False, False, False, False, False}
```

```
In[30]:= Table[With[{i = i}, ParallelEvaluate[IntegerQ[i], kernel]],  
 {i, 1, 10}]
```

```
Out[30]=
```

```
{True, True, True, True, True, True, True, True, True, True}
```

---

## ParallelCombine

```
ParallelCombine[f, h[e1, e2, ..., ei, ..., en], comb]
```

```
In[31]:= Prime[3]
```

```
Out[31]=
```

```
5
```

```
In[32]:= Prime[{1, 2, 3}]
```

```
Out[32]=
```

```
{2, 3, 5}
```

```
In[33]:= ParallelCombine[Prime, {1, 2, 3, 4, 5, 6}]
```

```
Out[33]=
```

```
{2, 3, 5, 7, 11, 13}
```

```
In[34]:= ParallelCombine[f, {1, 2, 3, 4, 5, 6}, comb]
```

```
Out[34]=
```

```
comb[f[{1, 2}], f[{3, 4}], f[{5}], f[{6}]]
```

```
In[35]:= ParallelCombine[f, h[1, 2, 3, 4, 5, 6], comb]
```

```
Out[35]=
```

```
comb[f[h[1, 2]], f[h[3, 4]], f[h[5]], f[h[6]]]
```

## ParallelMap

```
In[36]:= ParallelMap[Sin, {0,  $\pi$ , 1.0}]
```

```
Out[36]=
```

```
{0, 0, 0.841471}
```

```
In[37]:= ParallelMap[FactorInteger,  $\frac{10^{\text{Range}[20, 30]} - 1}{9}$ ]
```

Out[37]=

```

{{{11, 1}, {41, 1}, {101, 1}, {271, 1}, {3541, 1},
 {9091, 1}, {27961, 1}}, {{3, 1}, {37, 1}, {43, 1},
 {239, 1}, {1933, 1}, {4649, 1}, {10838689, 1}},
 {{11, 2}, {23, 1}, {4093, 1}, {8779, 1}, {21649, 1},
 {513239, 1}}, {{1111111111111111111111111111, 1}},
 {{3, 1}, {7, 1}, {11, 1}, {13, 1}, {37, 1}, {73, 1},
 {101, 1}, {137, 1}, {9901, 1}, {99990001, 1}},
 {{41, 1}, {271, 1}, {21401, 1}, {25601, 1},
 {182521213001, 1}}, {{11, 1}, {53, 1}, {79, 1},
 {859, 1}, {265371653, 1}, {1058313049, 1}}, {{3, 3},
 {37, 1}, {757, 1}, {333667, 1}, {440334654777631, 1}},
 {{11, 1}, {29, 1}, {101, 1}, {239, 1}, {281, 1},
 {4649, 1}, {909091, 1}, {121499449, 1}},
 {{3191, 1}, {16763, 1}, {43037, 1}, {62003, 1},
 {77843839397, 1}}, {{3, 1}, {7, 1}, {11, 1}, {13, 1},
 {31, 1}, {37, 1}, {41, 1}, {211, 1}, {241, 1},
 {271, 1}, {2161, 1}, {9091, 1}, {2906161, 1}}}

```

```
In[38]:= CloseKernels [ ]
```

```
Out[38]=
```

```
{KernelObject[9, local, <defunct>],  
 KernelObject[10, local, <defunct>]}
```



## Úlohy ParallelSubmit, WaitAll, ...

```
In[39]:= LaunchKernels [{"localhost", "localhost"}]
```

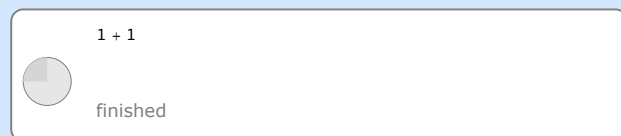
Out[39]=

```
{KernelObject[11, local], KernelObject[12, local]}
```

### Vytvoření a vyhodnocení úlohy

```
In[40]:= j = ParallelSubmit[1 + 1]
```

Out[40]=



1 + 1  
finished

```
In[41]:= WaitAll[j]
```

```
Out[41]=
```

```
2
```

```
In[42]:= pids = Function[i, ParallelSubmit[i^2]] /@ {1, 2, 3, 4, 5}
```

```
Out[42]=
```




```
In[43]:= {res, pid, pids} = WaitNext[pids]
```

Out[43]=



```
In[44]:= {res, pid, pids} = WaitNext[pids]
```

Out[44]=



```
In[45]:= WaitAll [pids]
```

```
Out[45]=
```

```
{9, 16, 25}
```

---

## Proměnné

```
In[46]:= a = 2
```

```
Out[46]=
```

```
2
```

```
In[47]:= Head [a]
```

```
Out[47]=
```

```
Integer
```

```
In[48]:= WaitAll [ParallelSubmit [Head [a] ] ]
```

```
Out[48]=
```

```
Symbol
```

```
In[49]:= With [ {a = a} , WaitAll [ParallelSubmit [Head [a] ] ] ]
```

```
Out[49]=
```

```
Integer
```

```
In[50]:= WaitAll [ParallelSubmit [ {a} , Head [a] ] ]
```

```
Out[50]=
```

```
Integer
```

- Table[]

```
In[51]:= Table [i2 , {i , 1 , 10} ]
```

```
Out[51]=
```

```
{1 , 4 , 9 , 16 , 25 , 36 , 49 , 64 , 81 , 100}
```

```
In[52]:= WaitAll [Table [ParallelSubmit [ $i^2$ ], { $i$ , 1, 10}]]]
```

Out[52]=

```
{ $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ ,  $i^2$ }
```

```
In[53]:= WaitAll [Table [ParallelSubmit [{ $i$ },  $i^2$ ], { $i$ , 1, 10}]]]
```

Out[53]=

```
{1, 4, 9, 16, 25, 36, 49, 64, 81, 100}
```

```
In[54]:= ParallelTable [ $i^2$ , { $i$ , 1, 10}]
```

Out[54]=

```
{1, 4, 9, 16, 25, 36, 49, 64, 81, 100}
```

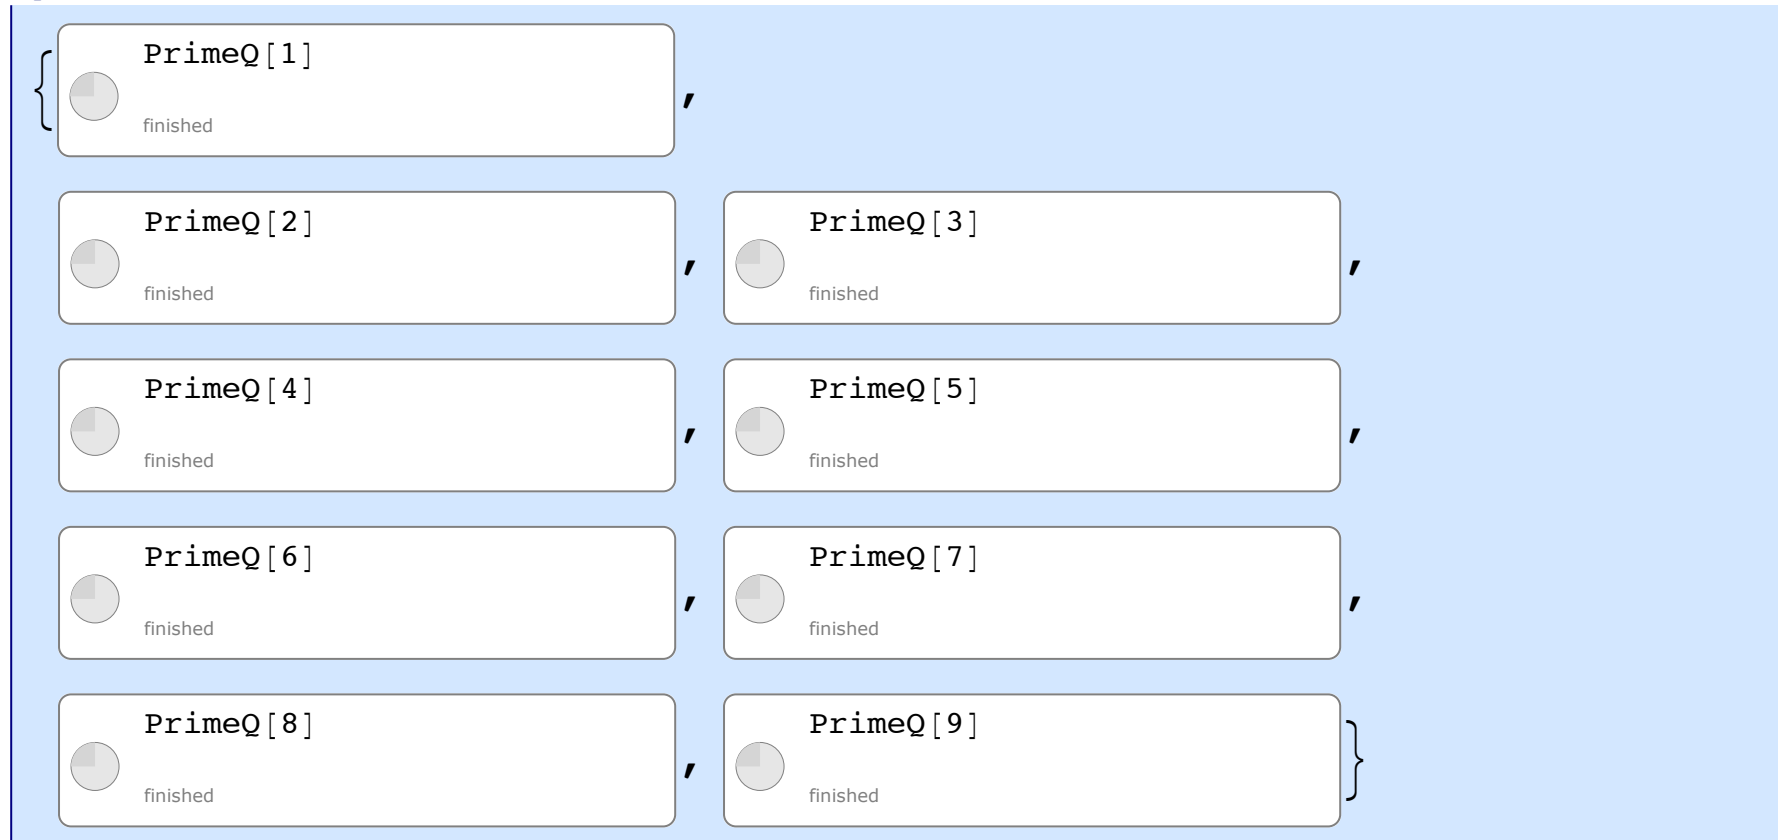
---

## Fronta úloh

```
In[55]:= Needs ["Parallel`Developer`"]
```

```
In[56]:= jobs = Table[ParallelSubmit[{i}, PrimeQ[i]], {i, 9}]
```

Out[56]=



```
In[57]:= QueueRun [ ]
```

```
Out[57]=
```

```
True
```

```
In[58]:= QueueRun [ ]
```

```
Out[58]=
```

```
True
```

```
In[59]:= $QueueLength
```

```
Out[59]=
```

```
5
```

```
In[60]:= WaitAll [jobs]
```

```
Out[60]=
```

```
{False, True, True, False, True, False, True, False, False}
```

```
In[61]:= CloseKernels [ ]
```

```
Out[61]=
```

```
{KernelObject[11, local, <defunct>],  
 KernelObject[12, local, <defunct>]}
```



## Distribuované definice `DistributeDefinitions[]` ...

```
In[62]:= LaunchKernels [{"localhost", "localhost"}]
```

Out[62]=

```
{KernelObject[13, local], KernelObject[14, local]}
```

### Použití proměnných a funkcí v paralelním prostředí

```
In[63]:= x = 10
```

Out[63]=

```
10
```

```
In[64]:= f[x_] := x^2
```

- Pozor, ačkoliv následující příklad funguje, ve skutečnosti nejde o paralelní vyhodnocení... (*proměnná ani funkce není definována na výpočetních jádrech, výpočet tak probíhá až na řídicím jádře*)

```
In[65]:= ParallelEvaluate[f[x]]
```

```
Out[65]=
```

```
{100, 100}
```

```
In[66]:= Head[x]
```

```
Out[66]=
```

```
Integer
```

```
In[67]:= ParallelEvaluate[Head[x]]
```

```
Out[67]=
```

```
{Symbol, Symbol}
```

```
In[68]:= DistributeDefinitions[x]
```

```
In[69]:= ParallelEvaluate[Head[x]]
```

```
Out[69]=
```

```
{Integer, Integer}
```

```
In[70]:= Clear[x]
```

```
In[71]:= Head[x]
```

```
Out[71]=
```

```
Symbol
```

```
In[72]:= ParallelEvaluate[Head[x]]
```

```
Out[72]=
```

```
{Integer, Integer}
```

```
In[73]:= DistributeDefinitions[x]
```

```
In[74]:= ParallelEvaluate[Head[x]]
```

```
Out[74]=
```

```
{Symbol, Symbol}
```

---

## Sdílené proměnné

```
In[75]:= x = 10;  
SetSharedVariable[x]
```

```
In[77]:= {k1, k2} = Kernels []
```

```
Out[77]=
```

```
{KernelObject[13, local], KernelObject[14, local]}
```

```
In[78]:= ParallelEvaluate[x, k1]
```

```
Out[78]=
```

```
10
```

```
In[79]:= ParallelEvaluate[x, k2]
```

```
Out[79]=
```

```
10
```

```
In[80]:= ParallelEvaluate[x = 42, k2]
```

```
Out[80]=
```

```
42
```

```
In[81]:= x
```

```
Out[81]=
```

```
42
```

```
In[82]:= ParallelEvaluate[x, k1]
```

```
Out[82]=
```

```
42
```

## Synchronizace - kritické sekce

```
In[83]:= SetSharedVariable[y];
```

- Sekvenční verze

```
In[84]:= y = 0;  
Map[ (  
  Pause[0.3 Random[]];  
  (* begin critical section *)  
  a = y;  
  Pause[Random[]];  
  y = a + 1  
  (* end critical section *)  
) &, Range[10]]
```

```
Out[85]=
```

```
{1, 2, 3, 4, 5, 6, 7, 8, 9, 10}
```

```
In[86]:= y
```

```
Out[86]=
```

```
10
```

- Paralelní verze

```
In[87]:= y = 0;  
ParallelMap[ (  
  Pause[0.3 Random[]];  
  (* begin critical section *)  
  a = y;  
  Pause[Random[]];  
  y = a + 1  
  (* end critical section *)  
  ) &,  
  Range[10] ]
```

```
Out[88]=
```

```
{1, 2, 3, 1, 2, 3, 4, 5, 6, 7}
```

```
In[89]:= y
```

```
Out[89]=
```

```
7
```

- Paralelní verze - ošetření kritické sekce - `CriticalSection[...]`

```
In[90]:= y = 0;  
ParallelMap[ (  
    Pause[0.3 Random[]];  
    CriticalSection[{lck},  
        a = y;  
        Pause[Random[]];  
        y = a + 1;  
        a + 1]) &,  
Range[10]]
```

```
Out[91]=
```

```
{2, 4, 6, 1, 3, 5, 7, 9, 8, 10}
```

```
In[92]:= y
```

```
Out[92]=
```

```
10
```

```
In[93]:= CloseKernels []
```

```
Out[93]=
```

```
{KernelObject[13, local, <defunct>],  
 KernelObject[14, local, <defunct>]}
```



# Lightweight Grid

## Konfigurace a spuštění výpočetních jader

Menu ▸ Evaluation ▸ Parallel Kernel Configuration...

In[94]:= **\$ConfiguredKernels**

```
{LightweightGridClient`LightweightGrid[{Agent →
  http://stroj1.domena:3737/WolframLightweightGrid/
  Manager, KernelCount → 4,
  LocalLinkMode → Connect, Service → , Timeout → 5}],
LightweightGridClient`LightweightGrid[
  {Agent →
  http://stroj2.domena:3737/WolframLightweightGrid/
  Manager, KernelCount → 8,
  LocalLinkMode → Connect, Service → , Timeout → 5}],
<<2 local kernels>> }
```

```
In[95]:= Needs [ "LightweightGridClient` " ]
```

```
In[96]:= Kernels [ ]
```

```
Out[96]=
```

```
{ }
```

```
LaunchKernels [ LightweightGrid [ "192.168.1.42" ] ]
```

```
Out[97]=
```

```
{ KernelObject [ 15, brian ] }
```

```
In[98]:= CloseKernels [ ]
```

```
Out[98]=
```

```
{ KernelObject [ 15, brian, <defunct> ] }
```