



Parallelism : The challenge for mainstream developers

Bringing Parallel Computing to the mass

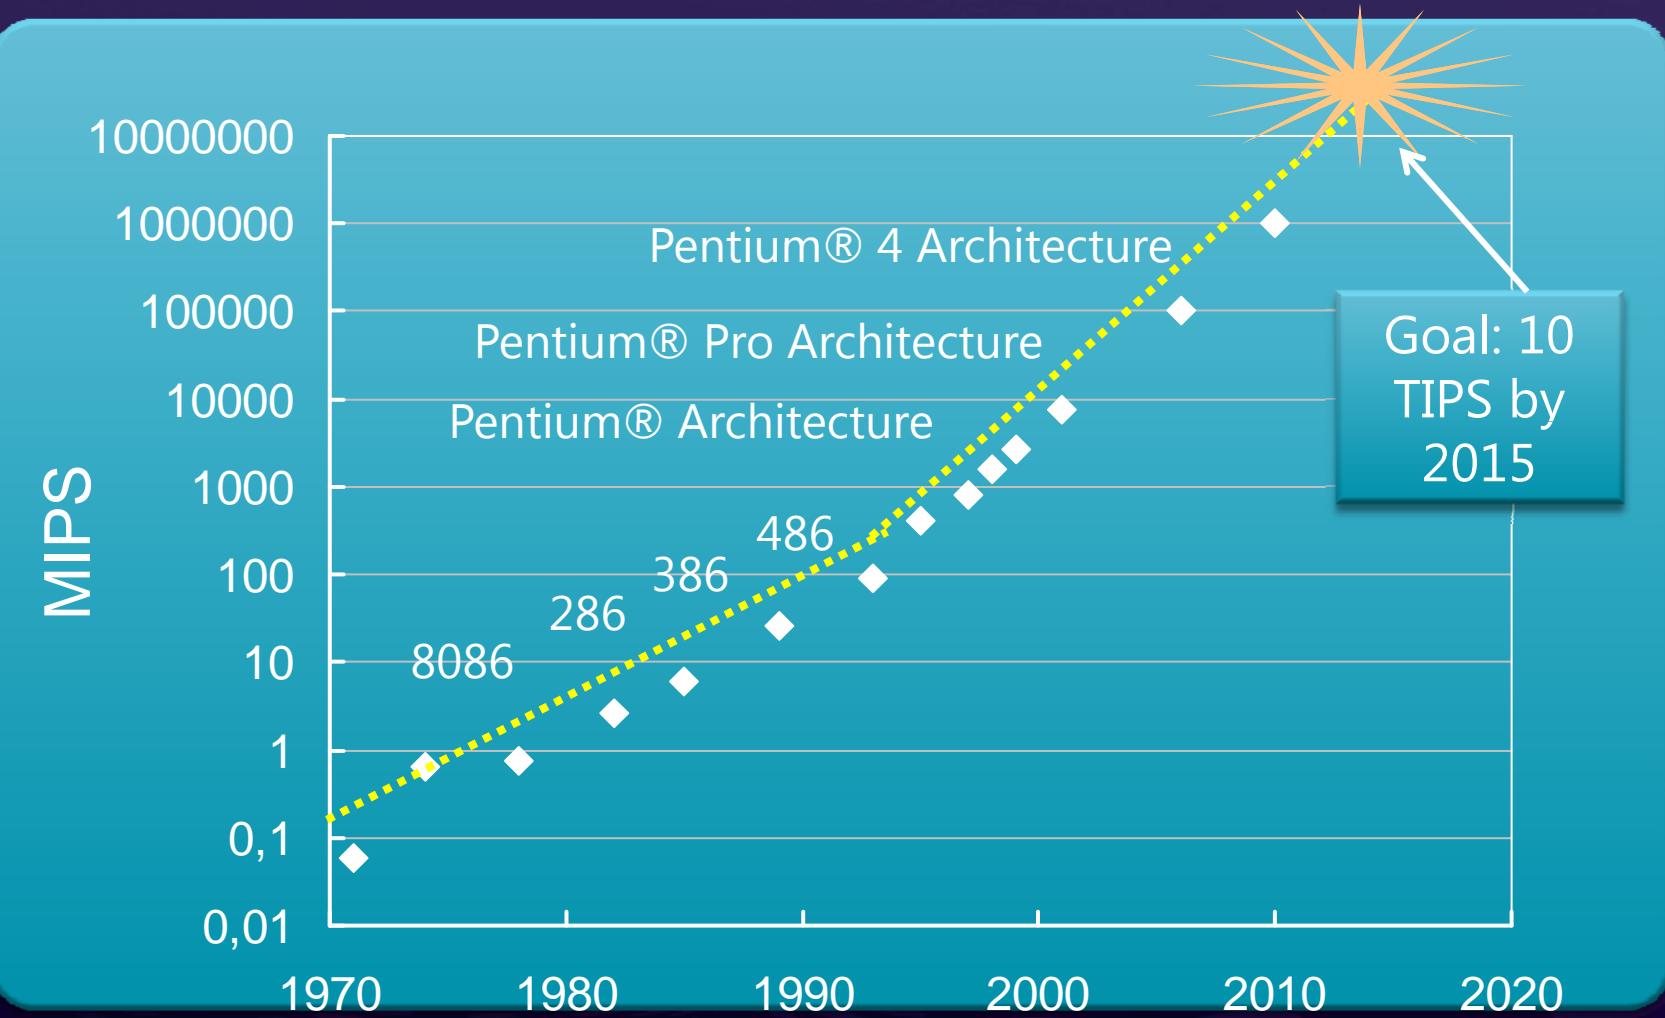
*Eric Vernié
Relation technique développeur
Microsoft France*



Gartner : 7 grand challenges

- Never having to manually recharge devices
- Parallel Programming
- Non Tactile, Natural Computing Interface
- Automated Speech Translation
- Persistent and Reliable Long-Term Storage
- Increase Programmer Productivity 100-fold
- Identifying the Financial Consequences of IT Investing

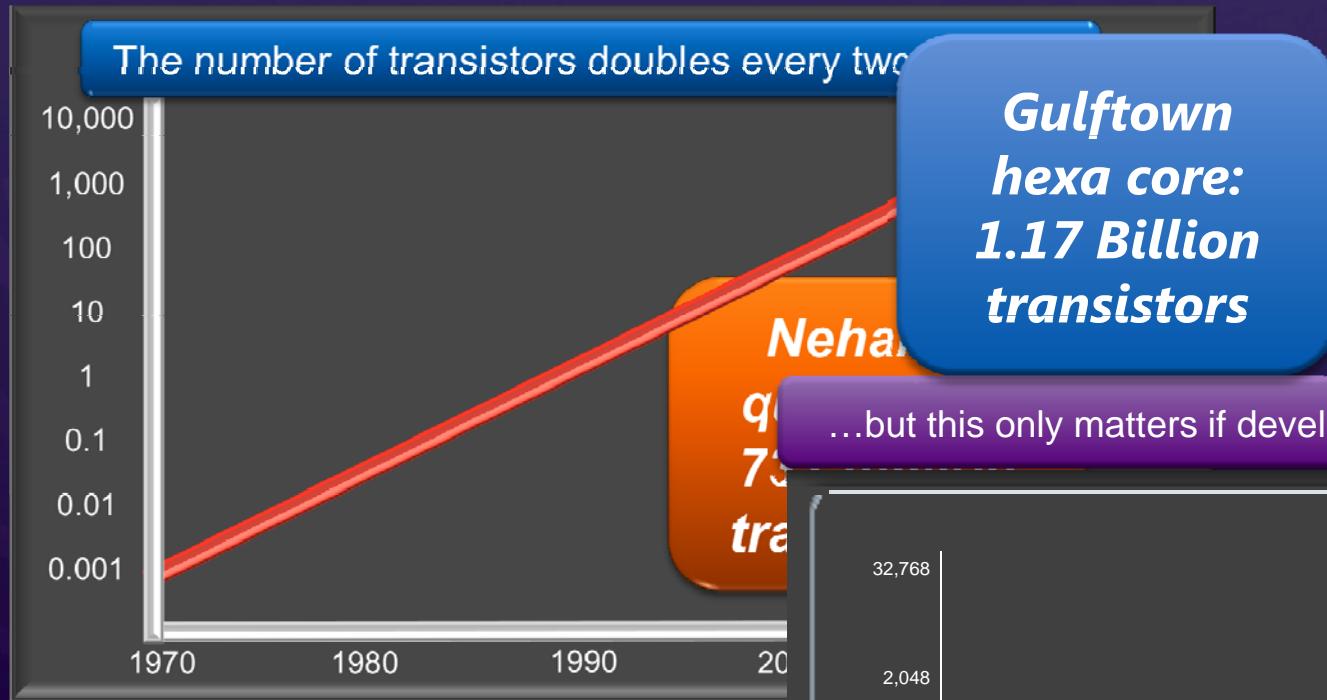
The physical laws change the rules



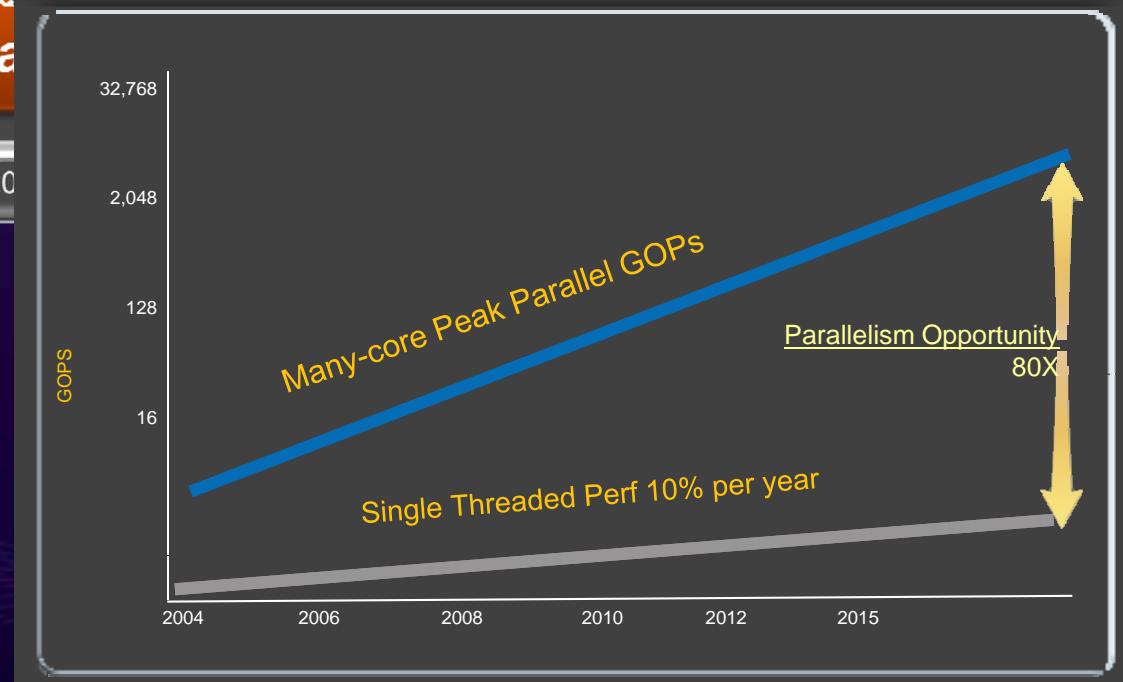
Source: Intel Juin 2006

Forum TER@TEC 2010 – 16 juin

The Manycore Shift



...but this only matters if developers embrace parallel computing!



In late 2008, only 3% of developers introduced parallelism in their applications

Evans Data

for mainstream developers, a new
story begin...

The manycore shift is a ***disruption*** for the developers

Free lunch is over

*If you haven't done so already, **now is the time** to take a hard look at the design of your application, determine what operations are CPU-sensitive now or are likely to become so soon, and **identify how those places could benefit from concurrency**."*

Herb Sutter, C++ Architect at Microsoft (March 2005)

What's the problem?

- Multithreaded programming is “**hard**” today
 - Doable by specialists
 - Parallel processing is known, nor easy to implement
 - So many ways to do it
- Businesses need “**simple steps**”
 - Best devs are not concurrency experts
 - Need simple ways to allow all devs to write concurrent code



So this is the challenge that we are focus on with our partner .

"Enable people to do parallelism without the need to go deep and make parallelism accessible to a large number of software developers."

Steve Texeira
Product Unit Manager
Parallel Developer Tools
Microsoft Corporation

Microsoft® Parallel Computing Initiative

**Microsoft's Parallel Computing Initiative encompasses
the vision, strategy, and innovative technologies
for delivering
transformative, natural, and immersive
personal computing experiences,
that harness the computing power of manycore architectures.**

The manycore is a ***disruption*** for the developer but it's also ***a great opportunity***

It begins with DEVELOPERS

Make it easier to express and manage the correctness, efficiency and maintainability of parallelism on Microsoft platforms for developers of all skill levels

Enable developers to express parallelism easily and focus on the problem to be solved

Improve the efficiency and scalability of parallel applications

Simplify the process of designing and testing parallel applications

Visual Studio 2010

Tools, Programming Models, Runtimes

Tools

Parallel Debugger

Visual Studio IDE

MPI Debugger

Programming Models

Parallel LINQ

Task Parallel Library

.NET Framework 4

ThreadPool

Task Scheduler

Resource Manager

Parallel Pattern Library

Agents Library

Visual C++ 10

Concurrency Runtime

Task Scheduler

Resource Manager

Operating System

Windows

UMS Threads

Key:

Managed

Native

Tooling

Tasks in .NET and C++

.NET 4.0

- Parallel.For(x, y, λ)
- Parallel.ForEach(IEnumerable, λ)
- Parallel.Invoke(λ, λ)
- Task
- Task.Factory.StartNew(λ)
- ThreadPool-based

Visual C++ 10

- parallel_for(x, y, step, λ);
- parallel_for_each(it, λ)
- parallel_invoke(λ, λ)
- task_group / task_handle
- task_group::run (λ)
- Native concurrency runtime

(and many overloads for the above)

Matrice Multiplication

```
void MatrixMul( int size, double** m1, double** m2, double** result)
{
    for (int i = 0; i < size; i++) {
        for (int j = 0; j < size; j++) {
            result[i][j] = 0;
            for (int k = 0; k < size; k++) {
                result[i][j] += m1[i][k] * m2[k][j];
            }
        }
    }
}
```

Manual solution

```
void MatrixMulThreading(int size, double** m1, double** m2, double** result) {
    int N = size;
    int P = size;
    int Q = size;
    HANDLE hEvent = CreateEvent(NULL, TRUE, FALSE, NULL);
    long count = 0;
    for (int c = 0; c < N; c++) {
        std::thread t[Q];
        for (int i = 0; i < N; i++) {
            for (int j = 0; j < P; j++) {
                result[i][j] = 0;
                for (int k = 0; k < Q; k++) {
                    result[i][j] += m1[i][k] * m2[k][j];
                }
            }
        }
        if (InterlockedCompareExchange((long*)&count, count + 1, 0) == 0)
            SetEvent(hEvent);
    }
    WaitForSingleObject(hEvent, INFINITE);
    CloseHandle(hEvent);
}
```

*"Nontrivial Software written with threads,
semaphores, and mutexes are incomprehensible
to humans and cannot and should not be trusted"*

Edward A. Lee
UC Berkeley

```
void MatrixMul( int size, double** m1, double** m2, double** result)
{
    parallel_for(0, size, 1, [&](int i) {
        for (int j = 0; j < size; j++) {
            result[i][j] = 0;
            for (int k = 0; k < size; k++) {
                result[i][j] += m1[i][k] * m2[k][j];
            }
        }
    });
}
```

Example: “Race Car Drivers”

```
IEnumerable<RaceCarDriver> drivers = ...;  
var results = new List<RaceCarDriver>();  
foreach(var driver in drivers)  
{  
    if (driver.Name == queryName &&  
        driver.Wins.Count >= queryWinCount)  
    {  
        results.Add(driver);  
    }  
}  
results.Sort((b1, b2) =>  
    b1.Age.CompareTo(b2.Age));
```

Manual Parallel Solution

```
IEnumerable<RaceCarDriver> drivers = ...;
var results = new List<RaceCarDriver>();
int partitionsCount = Environment.ProcessorCount;
int remainingCount = partitionsCount;
var enumerator = drivers.GetEnumerator();
try {
    using (var done = new ManualResetEvent(false)) {
        for(int i = 0; i < partitionsCount; i++) {
            ThreadPool.QueueUserWorkItem(delegate {
                while(true) {
                    RaceCarDriver driver;
                    lock (enumerator) {
                        if (!enumerator.MoveNext()) break;
                        driver = enumerator.Current;
                    }
                    if (driver.Name == queryName &&
                        driver.Wins.Count >= queryWinCount) {
                        lock(results) results.Add(driver);
                    }
                }
                if (Interlocked.Decrement(ref remainingCount) == 0) done.Set();
            });
        }
        done.WaitOne();
        results.Sort((b1, b2) => b1.Age.CompareTo(b2.Age));
    }
}
finally { if (enumerator is IDisposable) ((IDisposable)enumerator).Dispose(); }
```

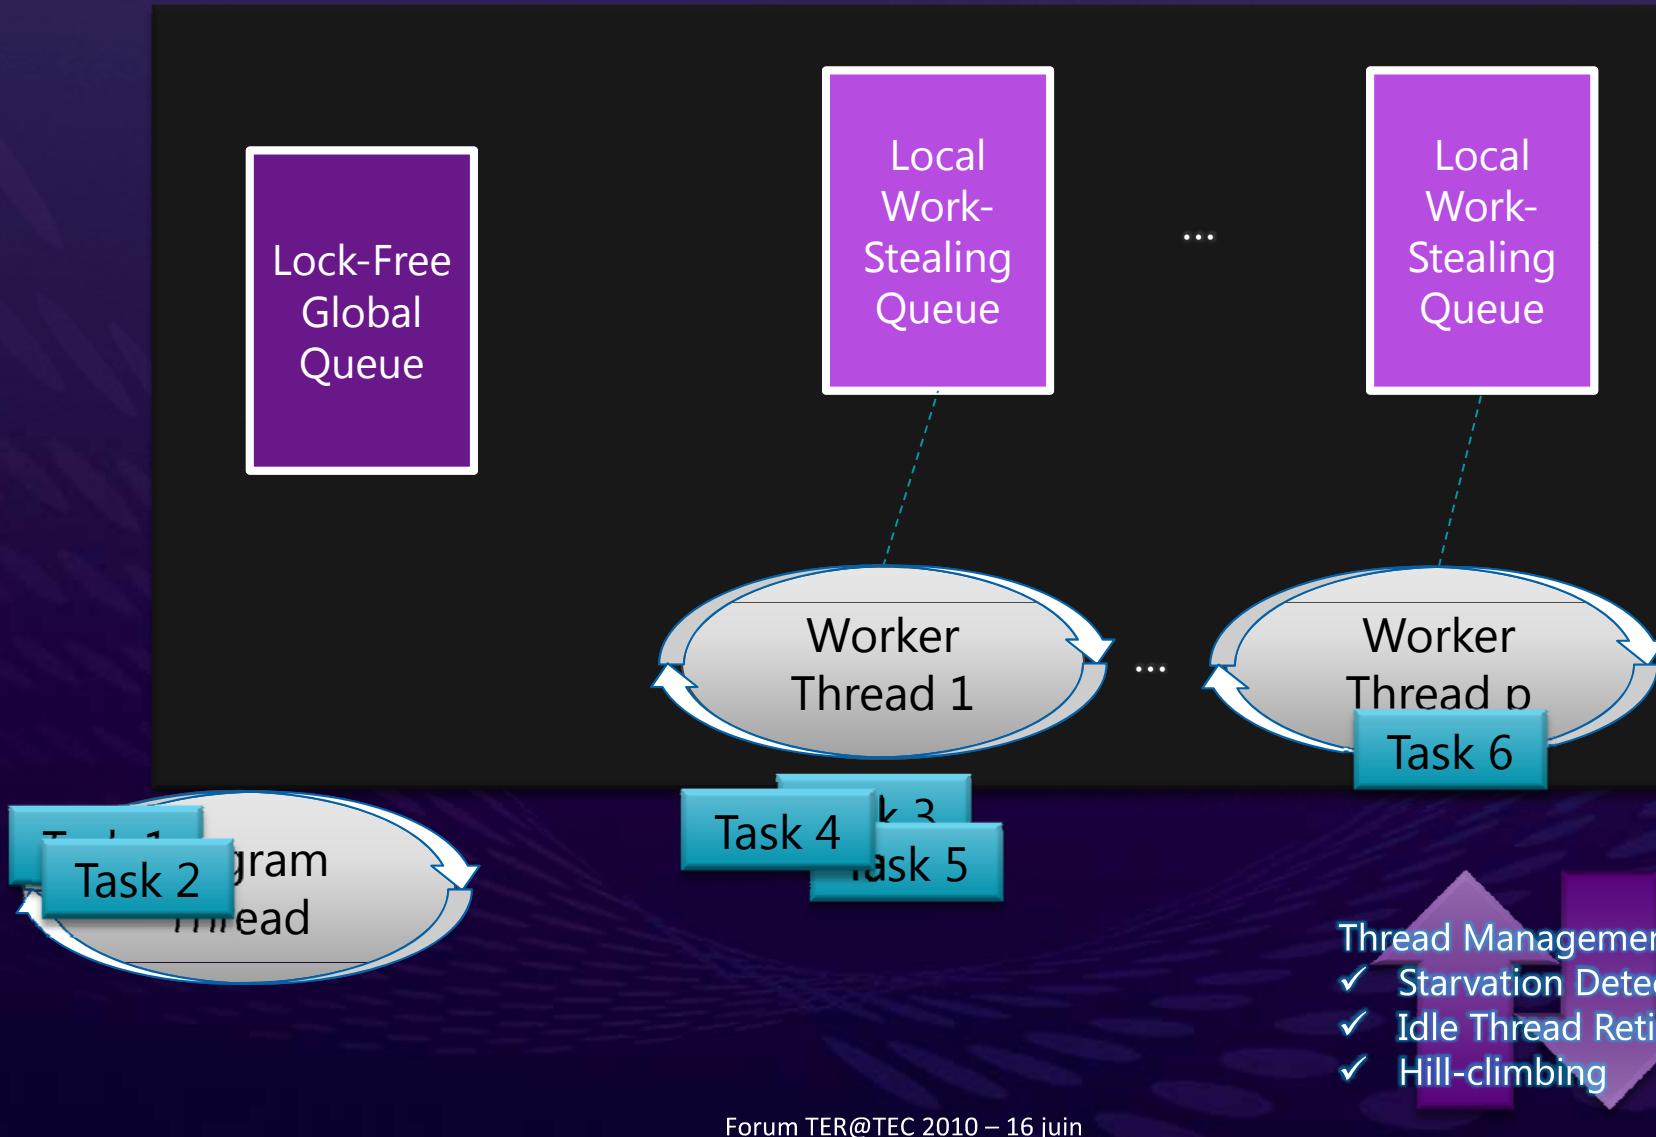
LINQ Solution

```
var results = from driver in drivers
    where driver.Name == queryName &&
        driver.Wins.Count >= queryWinsCount
    orderby driver.Age ascending
    select driver;
```

Concurrency Runtime

- Programming-Model-Independent
- Scheduler
 - Low-overhead scheduling and load balancing
 - Policy-driven
 - Multiple scheduler instances, cooperative blocking, supports UMS on Win7
 - APIs
 - Scheduler, SchedulerPolicy, ScheduleGroup, Context
- Resource Manager
 - Serves requests for physical processors, intra-process
 - Dynamic load balancing between scheduler instances
 - Interfaces
 - IResourceManager, IScheduler, IVirtualProcessorRoot, IThreadProxy, IExecutionContext

Work-Stealing Scheduler



Debugging & profiling tools

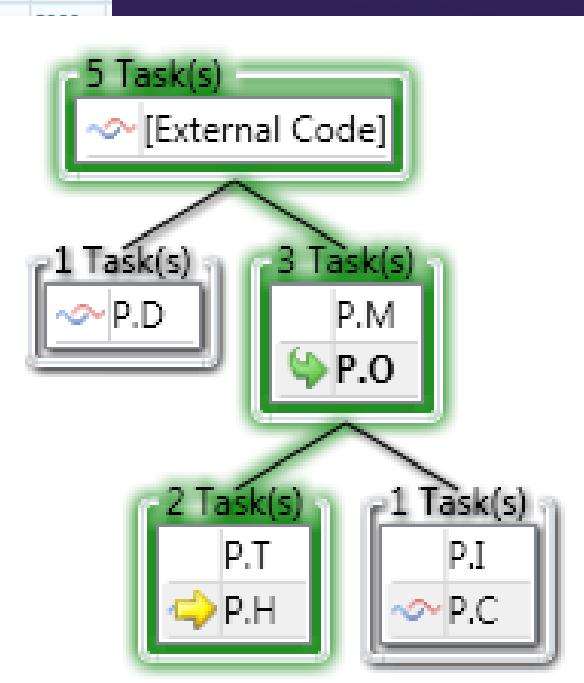
Tasks					
	Id.	Status	Location	Parent Task	Thread
P	▲ 1242864	Waiting	Concurrency::StructuredTaskPool::Wait		-----
P	▲ 1963792	Waiting	Concurrency::StructuredTaskPool::Wait	1242864	
P	▲ 1961324	Waiting	Concurrency::StructuredTaskPool::Wait	1963792	
P	▲ 1958136	Waiting	Concurrency::StructuredTaskPool::Wait	1961324	
P	1954948	Scheduled	0x004110be CChore::Payload(void *)	1958136	
P	1955128	Running	State::Remove	1958136	
P	▲ 1243044	Waiting	Concurrency::StructuredTaskPool::Wait		
P	▲ 1701972	Waiting	Concurrency::StructuredTaskPool::Wait	1243044	
P	▲ 1698964	Waiting	Concurrency::StructuredTaskPool::Wait	1701972	
P	1695776	Scheduled	0x004110be CChore::Payload(void *)	1698964	
P	▲ 1695956	Waiting	Concurrency::StructuredTaskPool::Wait	1698964	
P	1692108	Scheduled	0x004110be CChore::Payload(void *)	1695956	
P	1692288	Running	State::Solve	1695956	

Concurrency Visualization
Visualizes the behavior of a multi-threaded application


CPU Utilization


Threads


Cores



The diagram illustrates a task hierarchy. At the top is a box labeled "5 Task(s)" containing "[External Code]". Below it, two boxes branch out: one labeled "1 Task(s)" containing "P.D" and another labeled "3 Task(s)" containing "P.M" and "P.O". The "P.M" box further branches into "2 Task(s)" containing "P.T" and "P.H", and "1 Task(s)" containing "P.I" and "P.C". Each box contains a small icon representing its state: a wavy line for external code, a yellow arrow for P.D, a green arrow for P.M, a green circle for P.O, a yellow arrow for P.T, a yellow arrow for P.H, a wavy line for P.I, and a wavy line for P.C.

Baby Step

Thinking Parallel

*As thinking Parallel becomes intuitive,
structuring problems to scale will
become second nature*

James Reinder

Chief Software Evangelist and Director of Software Development
Products at Intel Corporation

question & answer