



HP's Hidden Supercomputers

Bill Dyck
Michael C Smith

Hewlett-Packard
Corporate-IT
3000 Hanover Street
Palo Alto, CA 94304
(916) 748-7248
bill_dyck@hp.com
michael_c_smith@hp.com

Peer-to-Peer
Computing at HP



Overview of Presentation



- The P2P Industry
- Definition of what we mean by P2PC
- Overview of P2PC Computing History
- Open Source Activities regarding P2P
- P2PC Business and Application Models
- Potential for P2PC in HP
 - Idle CPU measurements
 - Comparison using industry benchmarks
 - Millipede – Demonstration of P2PC implementation in HP

Peer-to-Peer Industry

- Peer-to-Peer File Sharing
- Peer-to-Peer Networking
- Peer-to-Peer Collaboration
- Peer-to-Peer Computing

What Is Peer-to-Peer Computing?

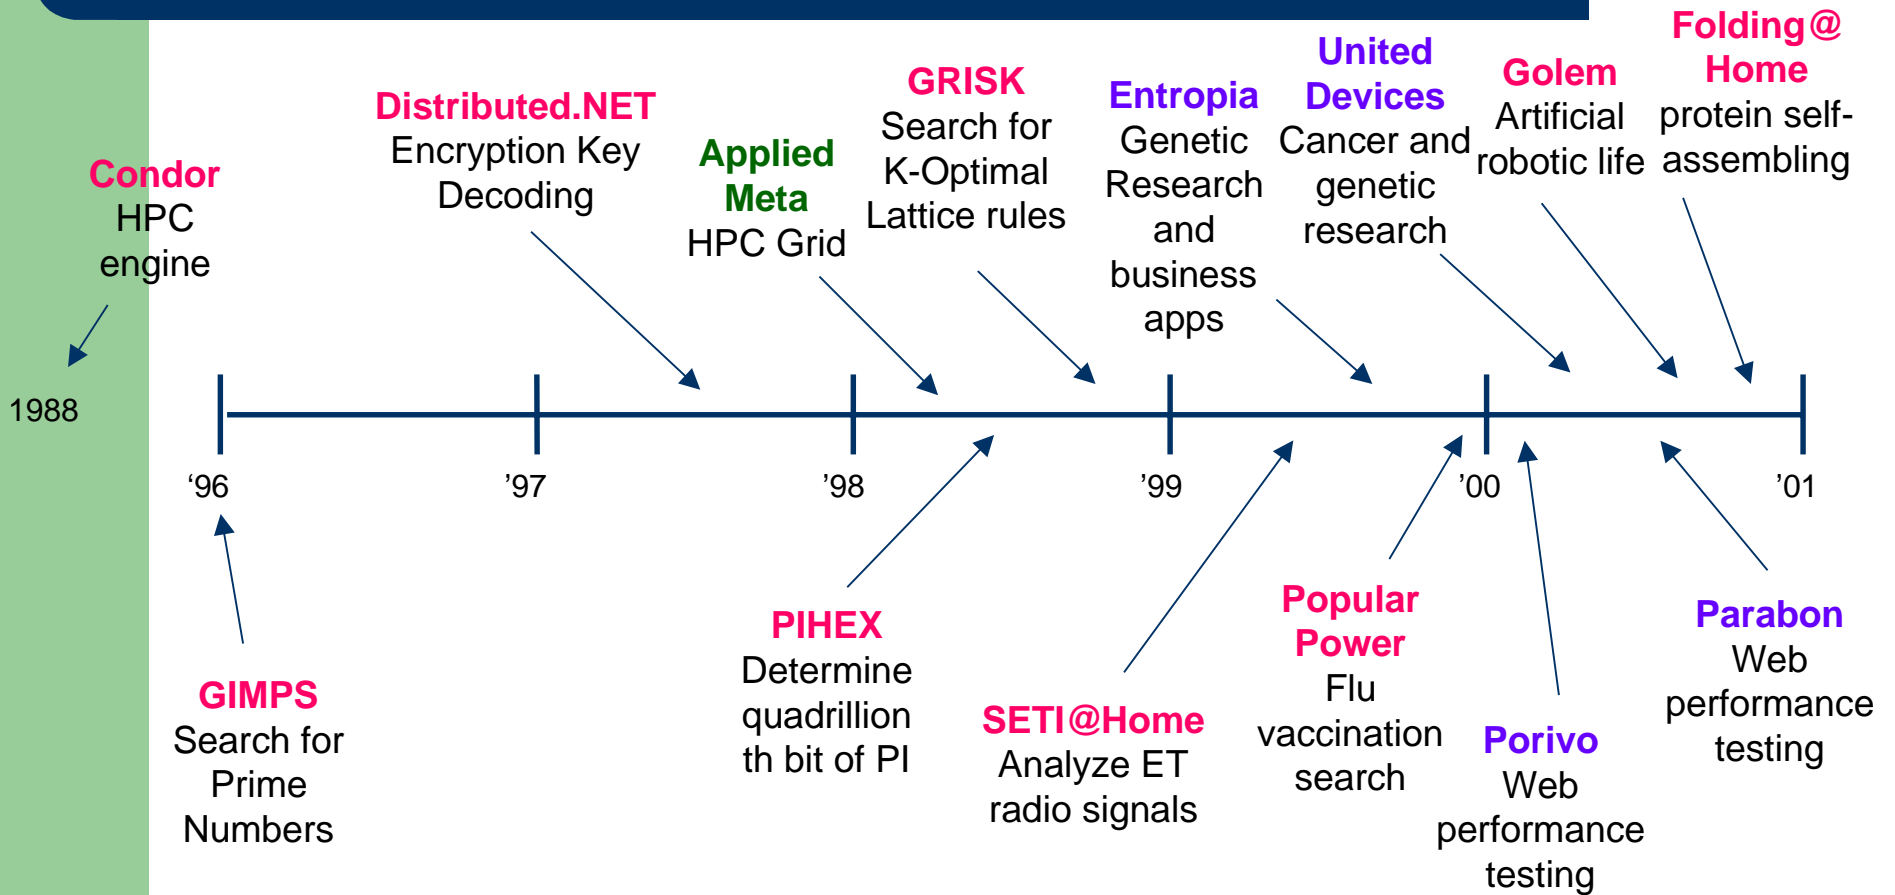
- Distributed Computing
 - Remote execution of applications
 - Remote computer is usually a dedicated server
- P2PC is a subset of Distributed Computing
 - Still remote execution of applications
 - Central server not responsible for processing power – that exists on other computers, usually on the fringe of the Internet cloud
 - Uses spare processing power



Legend

- Scientific
- Business
- Hybrid (Business/Scientific)

Overview of P2PC History



P2PC Open Source Activities

- Sun JXTA - www.jxta.org
 - Announced in February, 2001, beta released April, 2001
 - Applies to the whole P2P industry
 - Provides a set of “APIs” that simplify the development of all types of P2P applications
 - JXTA Objectives (from [Project JXTA: A Technology Overview](#))
 - Interoperability
 - Platform Independence
 - Ubiquity
 - Not specific to P2PC

P2PC Open Source Activities


- Intel P2P Working Group - www.p2pwg.org
 - Objective is to “facilitate and accelerate the advancement of infrastructure best-known practices for *peer-to-peer computing*.”
 - Wide industry participation, starting August, 2000.
 - Released a set of “API”s called PtPTL – Trusted Library
 - Provides security APIs
 - Provides simple network and OS support
 - Portable to both Win32 and Linux
 - Not specific to P2PC, even though the working group uses the term peer-to-peer computing

P2PC Business Models

- Philanthropic (mainly scientific)
- Provider (of CPU resources)
 - Direct payment to top CPU providers
 - Prizes, drawing, sweepstakes
 - Reverse Incentive
- Consumer (of CPU resources)
 - Grid of CPU cycles available for purchase
 - Can exist inside or outside the firewall

P2PC Application Models

- Centrally Managed, Central Data Collection



- Centrally Managed, Remote Data Collection



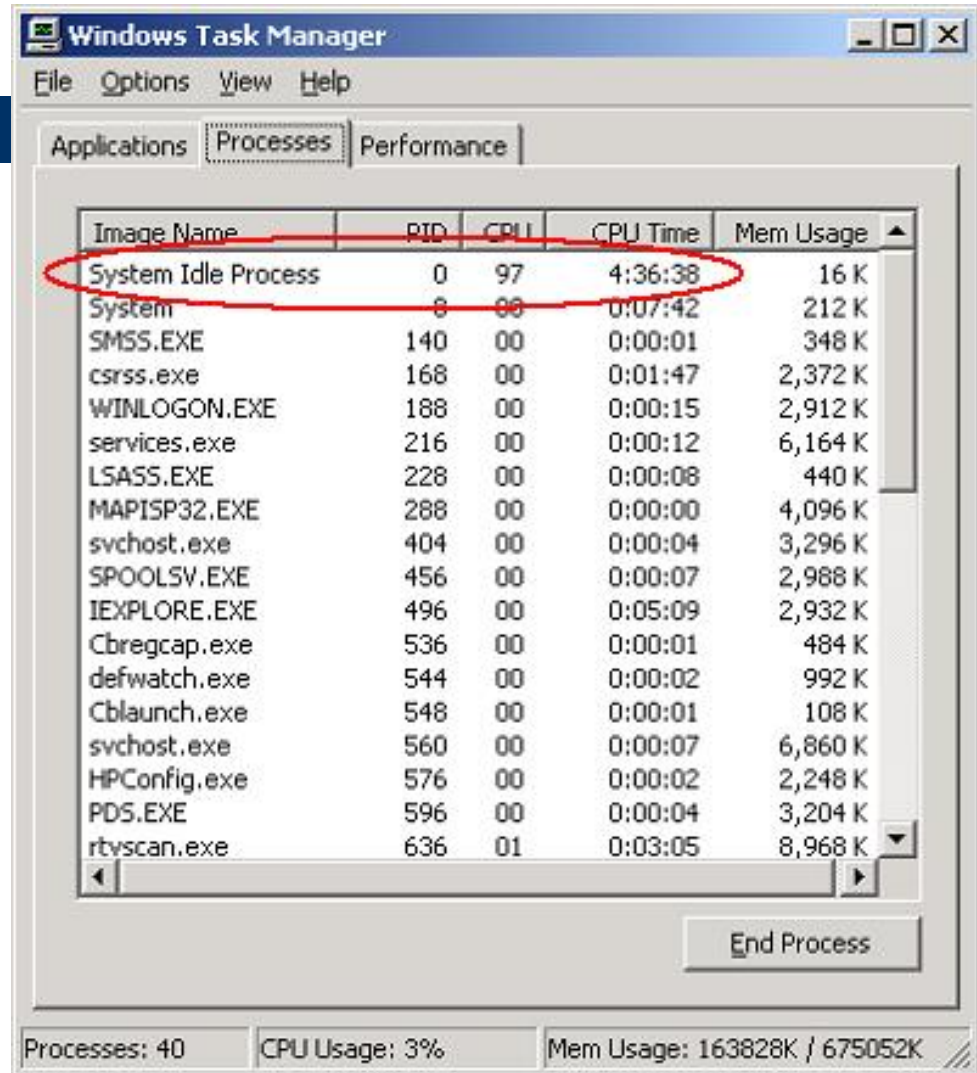
How to view your own spare CPU capacity

The System Idle Process is executed while your CPU is not busy doing something else.

Process

- Press CTRL-SHFT-ESC to bring up the Windows Task Manager.
- Click on the **Processes** tab.
- View the **System Idle Process** and its associated **CPU Time**.

Idle: 4 hours, 36 minutes



Windows Task Manager

File Options View Help

Applications Processes Performance

Image Name	PID	CPU	CPU Time	Mem Usage
System Idle Process	0	97	4:36:38	16 K
System	4	00	0:07:42	212 K
SMSS.EXE	140	00	0:00:01	348 K
csrss.exe	168	00	0:01:47	2,372 K
WINLOGON.EXE	188	00	0:00:15	2,912 K
services.exe	216	00	0:00:12	6,164 K
LSASS.EXE	228	00	0:00:08	440 K
MAPISP32.EXE	288	00	0:00:00	4,096 K
svchost.exe	404	00	0:00:04	3,296 K
SPOOLSV.EXE	456	00	0:00:07	2,988 K
IEXPLORE.EXE	496	00	0:05:09	2,932 K
Cbregcap.exe	536	00	0:00:01	484 K
defwatch.exe	544	00	0:00:02	992 K
Cblaunch.exe	548	00	0:00:01	108 K
svchost.exe	560	00	0:00:07	6,860 K
HPConfig.exe	576	00	0:00:02	2,248 K
PDS.EXE	596	00	0:00:04	3,204 K
rtvscan.exe	636	01	0:03:05	8,968 K

End Process

Processes: 40 CPU Usage: 3% Mem Usage: 163828K / 675052K


Overview of Idle CPU Measurements

- 88,167 CPU readings from 46 computers over a period of two weeks.
- Users are typical knowledge workers using MS products
- Readings occurred every 10 seconds between 8:00 am and 5:00 pm, starting when the user booted up their computer.
 - 7.93% Average CPU utilization overall
 - Standard deviation - 6.85%
 - Highest system average utilization – 29.37%
 - Lowest system average utilization – 0.25%

Details on Idle CPU Measurements


➤ Overall view of idle measurements

100% utilization represents a 100% busy CPU



➤ CPU utilization breakdown

- <1% CPU utilization accounts for 45% of ALL readings
- <10% CPU utilization accounts for 82% of ALL readings



Industry Benchmarks

- Standards Performance Evaluation Corporation (SPEC)
- Benchmarks (CPU2000, CPU95) - provide a comparative measure of compute intensive performance on various CPUs.
- Useful as a starting point to roughly determine available capacity among various CPUs.

Industry Benchmark (data)

Platform	int2000	fp2000	int95	fp95
166Mhz Pentium MMX			5.6	4.34
266Mhz Pentium Pro			7.73	6.57
433Mhz Celeron			14.9	10.9
450Mhz Pentium II			18.5	13.3
500Mhz Pentium II Xeon			21.6	15.9
650Mhz Pentium III	299	215		
1Ghz Pentium III	442	284		
1.7Ghz Pentium 4	586	608		
500Mhz SGI Origin3200 R14K	427	463		
750Mhz HP J6700 PA-8700	603	581		
900Mhz Sun UltraSPARC-III	467	482		

P2PC Comparison Formula

How can you compare one CPU performance to another?

$$\frac{1}{\text{perf}^{S1} / \text{perf}^{S2} * \%^{AV} * \%^{RC}} = \text{count}^{S1}$$

- count^{S1} = number of “S1” systems required to equal performance of “S2” system
- perf^{S1} = performance of S1 system
- perf^{S2} = performance of S2 system
- $\%^{AV}$ = percent of idle CPU time available on S1
- $\%^{RC}$ = percent of idle CPU time able to be harvested from S1

P2PC Comparison Example

Question: How many 166Mhz Pentium's would it take to equal the performance of a 1.7Ghz Pentium 4?

$$\frac{1}{\text{perf}^{S1} / \text{perf}^{S2} * \%^{AV} * \%^{RC}} = \text{count}^{S1}$$

- $\text{perf}^{S1} = 5.6$
- $\text{perf}^{S2} = 586$
- $\%^{AV} = 90\%$ available CPU cycles
- $\%^{RC} = 50\%$ harvestable CPU cycles

$$\frac{1}{5.6 / 586 * .9 * .5} = 232 \text{ Pentium 166's}$$



Millipede

A Demonstration of P2PC Implementation at HP

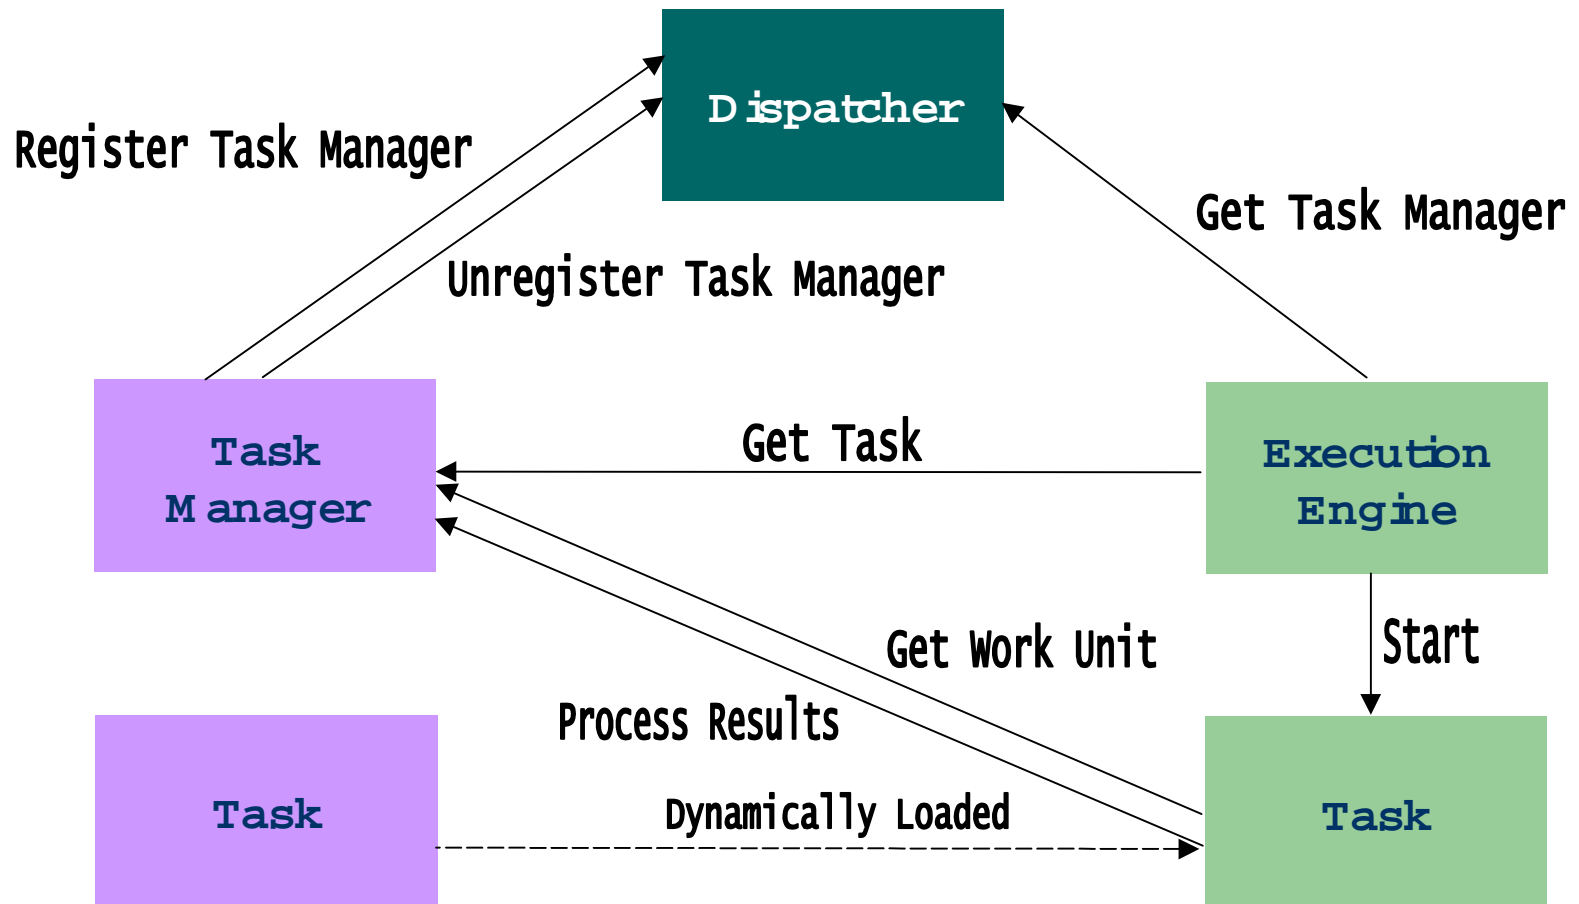
- Runs plug-in applications
- Real-time application deployment
- Policy-based security model
- Runs on any Java-enabled system

Millipede

Desirable Attributes

- Extensible
- Rapid deployment of applications
- Secure
- Heterogeneous
- Cost effectively runs plug-in applications

Millipede Architecture



Web References

Condor	http://www.cs.wisc.edu/condor	Parabon	http://www.parabon.com/
GIMPS	http://www.mersenne.org/prime.htm	Golem	http://www.demo.cs.brandeis.edu/golem
Distributed.NET	http://www.distributed.net/	Folding @ Home	http://www.stanford.edu/group/pandegrou/Cosm
Applied Meta	http://www.appliedmeta.com	Juno	http://www.juno.com
PIHEX	http://www.cecm.sfu.ca/projects/pihex	Sun's JXTA	http://www.jxta.org
GRISK	http://www.ii.uib.no/grisk	Intel P2P Working Group	http://www.p2pwg.org
SETI @ Home	http://setiathome.ssl.berkeley.edu	SPEC	http://www.spec.org
Entropia	http://www.entropia.com		
Popular Power	http://www.popularpower.com		
Porivo	http://www.porivo.com		
United Devices	http://www.uniteddevices.com		