Classification(s) (used in the following ...)

- Flynn (1972)
  - SISD, SIMD, MISD, MIMD
- MIMD
  - distributed memory, shared memory, distributed shared memory
- Memory access
  - UMA, NUMA, NORMA
- Interconnection networks
  - bus, crossbar, mesh, hypercube, fat-tree, ...

Moore’s Law

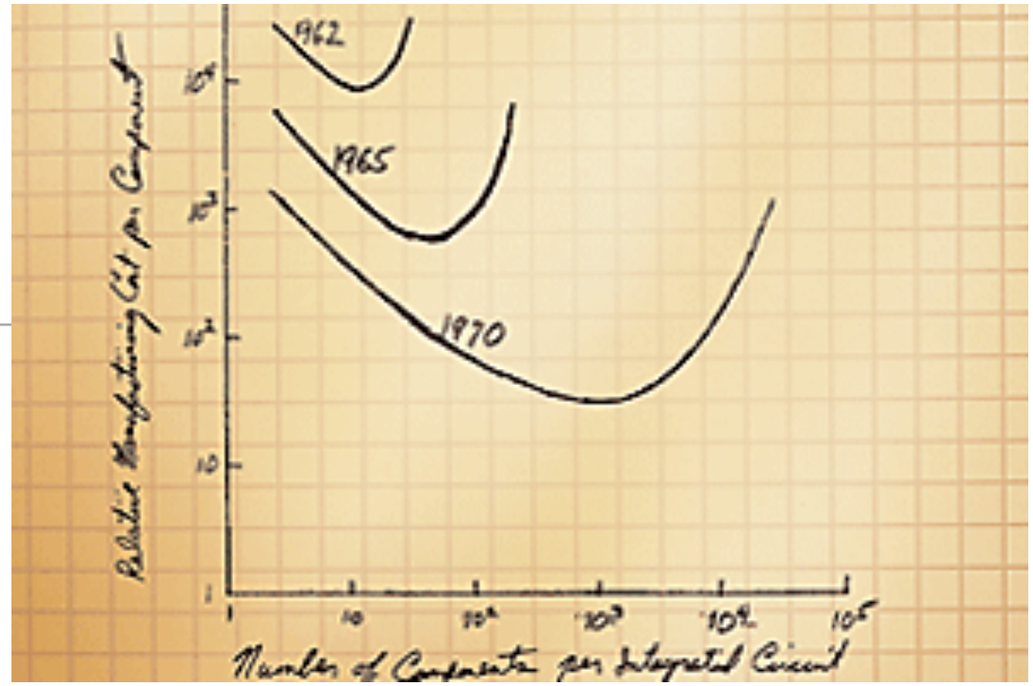

- Moore's original statement can be found in his publication "Cramming more components onto integrated circuits", Electronics Magazine 19 April 1965:

- “The complexity for minimum component costs has increased at a rate of roughly a factor of two per year ... Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years. That means by 1975, the number of components per integrated circuit for minimum cost will be 65,000. I believe that such a large circuit can be built on a single wafer.”

Moore - Intel

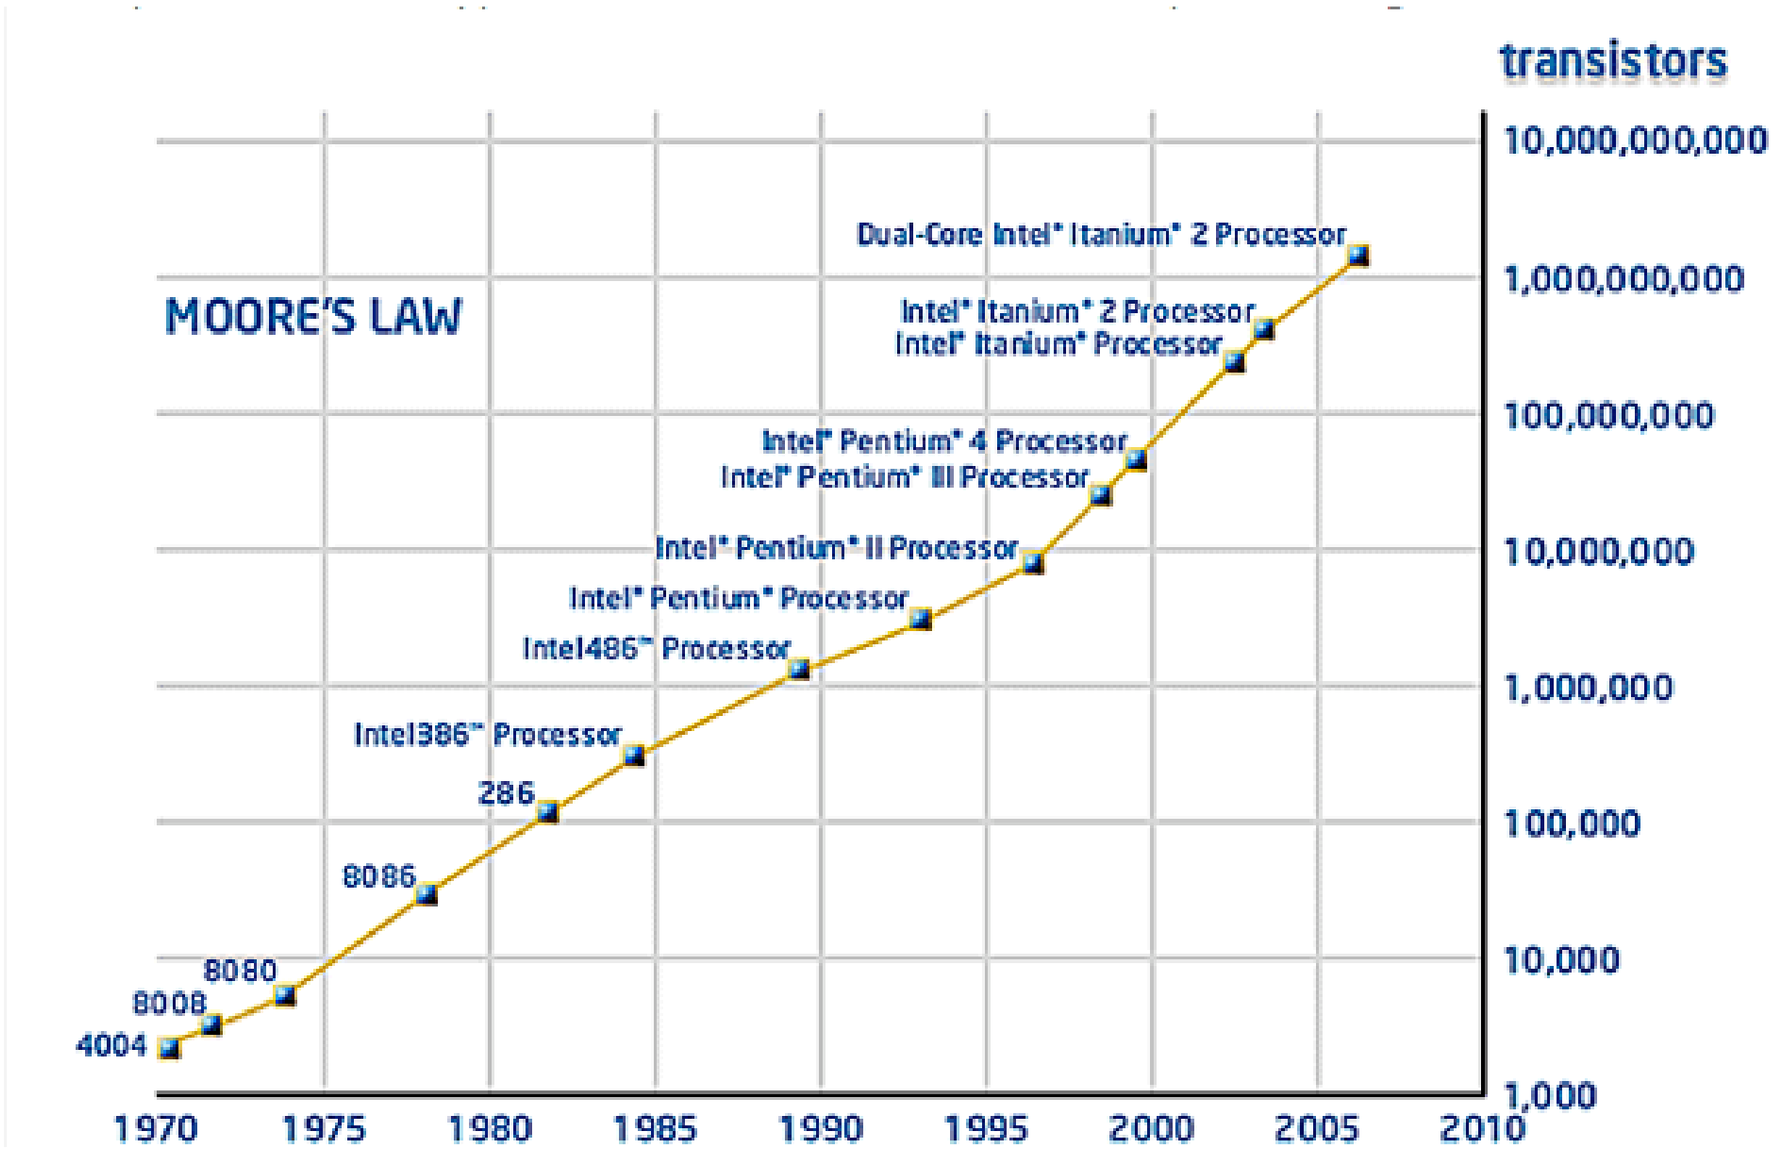
transistors
MOORE'S LAW
Dual-Core Intel® Itanium® 2 Processor
Intel® Itanium® 2 Processor
Intel® Itanium® Processor
Intel® Pentium® 4 Processor
Intel® Pentium® III Processor
Intel® Pentium® II Processor
Intel® Pentium® Processor
Intel486™ Processor
Intel386™ Processor
286
8086
8080
8008
4004
10,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
1,000
1970
1975
1980
1985
1990
1995
2000
2005
2010

Moore’s law (revisited)

- The original one:
  - number of active components/transistors doubles each year and a half
- limits:
  - usage ? processor model ?
  - big chip areas for small performance improvements !
- and therefore
  - law moved to cores per chip !!!

Top500

- rating publiches every 6 months (june and november)
  - at the same time of Supercomputing conference
- precise benchmark suite
- measures (raw) computing power
  - but also the trend (archive from 1993)

Evolution: architectures

- 1993

Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
SIMD	35	7.00 %	64	135	54272
MPP	119	23.80 %	400	826	14766
Single Processor	97	19.40 %	147	186	99
SMP	249	49.80 %	511	640	1983
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

- 2000

Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Constellations	93	18.60 %	6941	11187	10432
MPP	257	51.40 %	48026	72377	92081
Cluster	11	2.20 %	2260	3719	3668
SMP	139	27.80 %	7003	8392	9477
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

- 2006

Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Constellations	31	6.20 %	156728	250009	51274
MPP	108	21.60 %	1552019	1970903	484370
Cluster	361	72.20 %	1818711	2992636	485317
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

Evolution: architectures

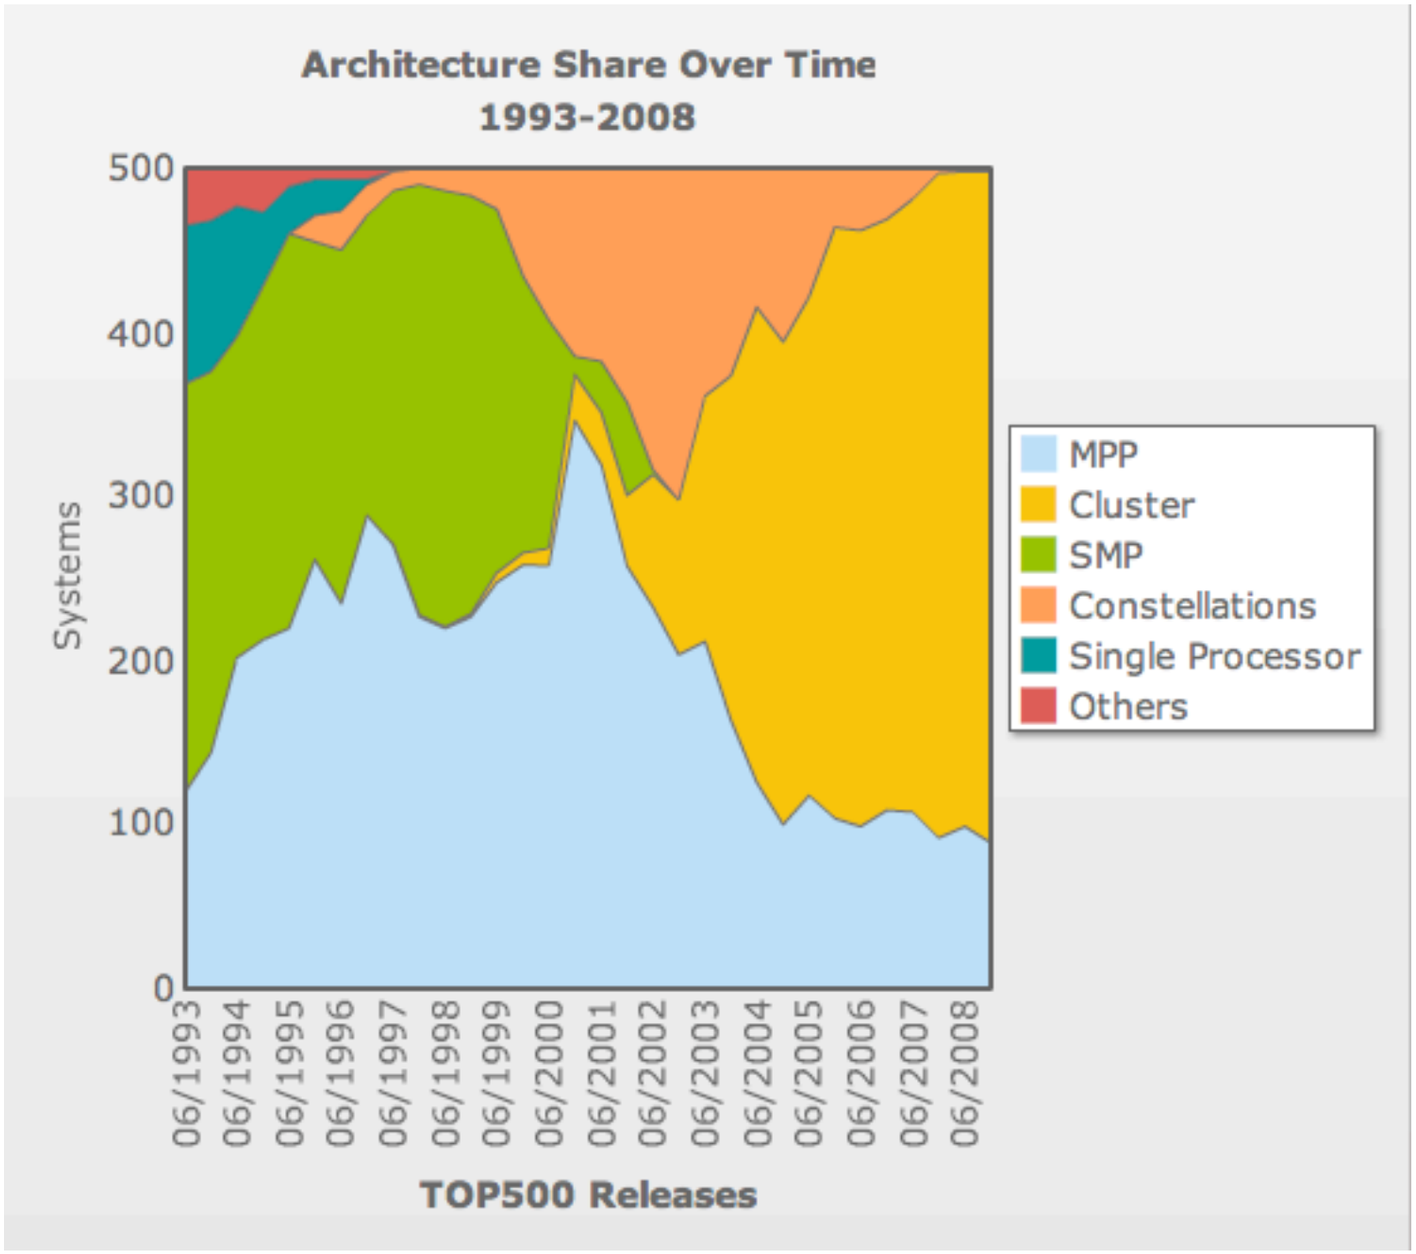

Evolution: architectures

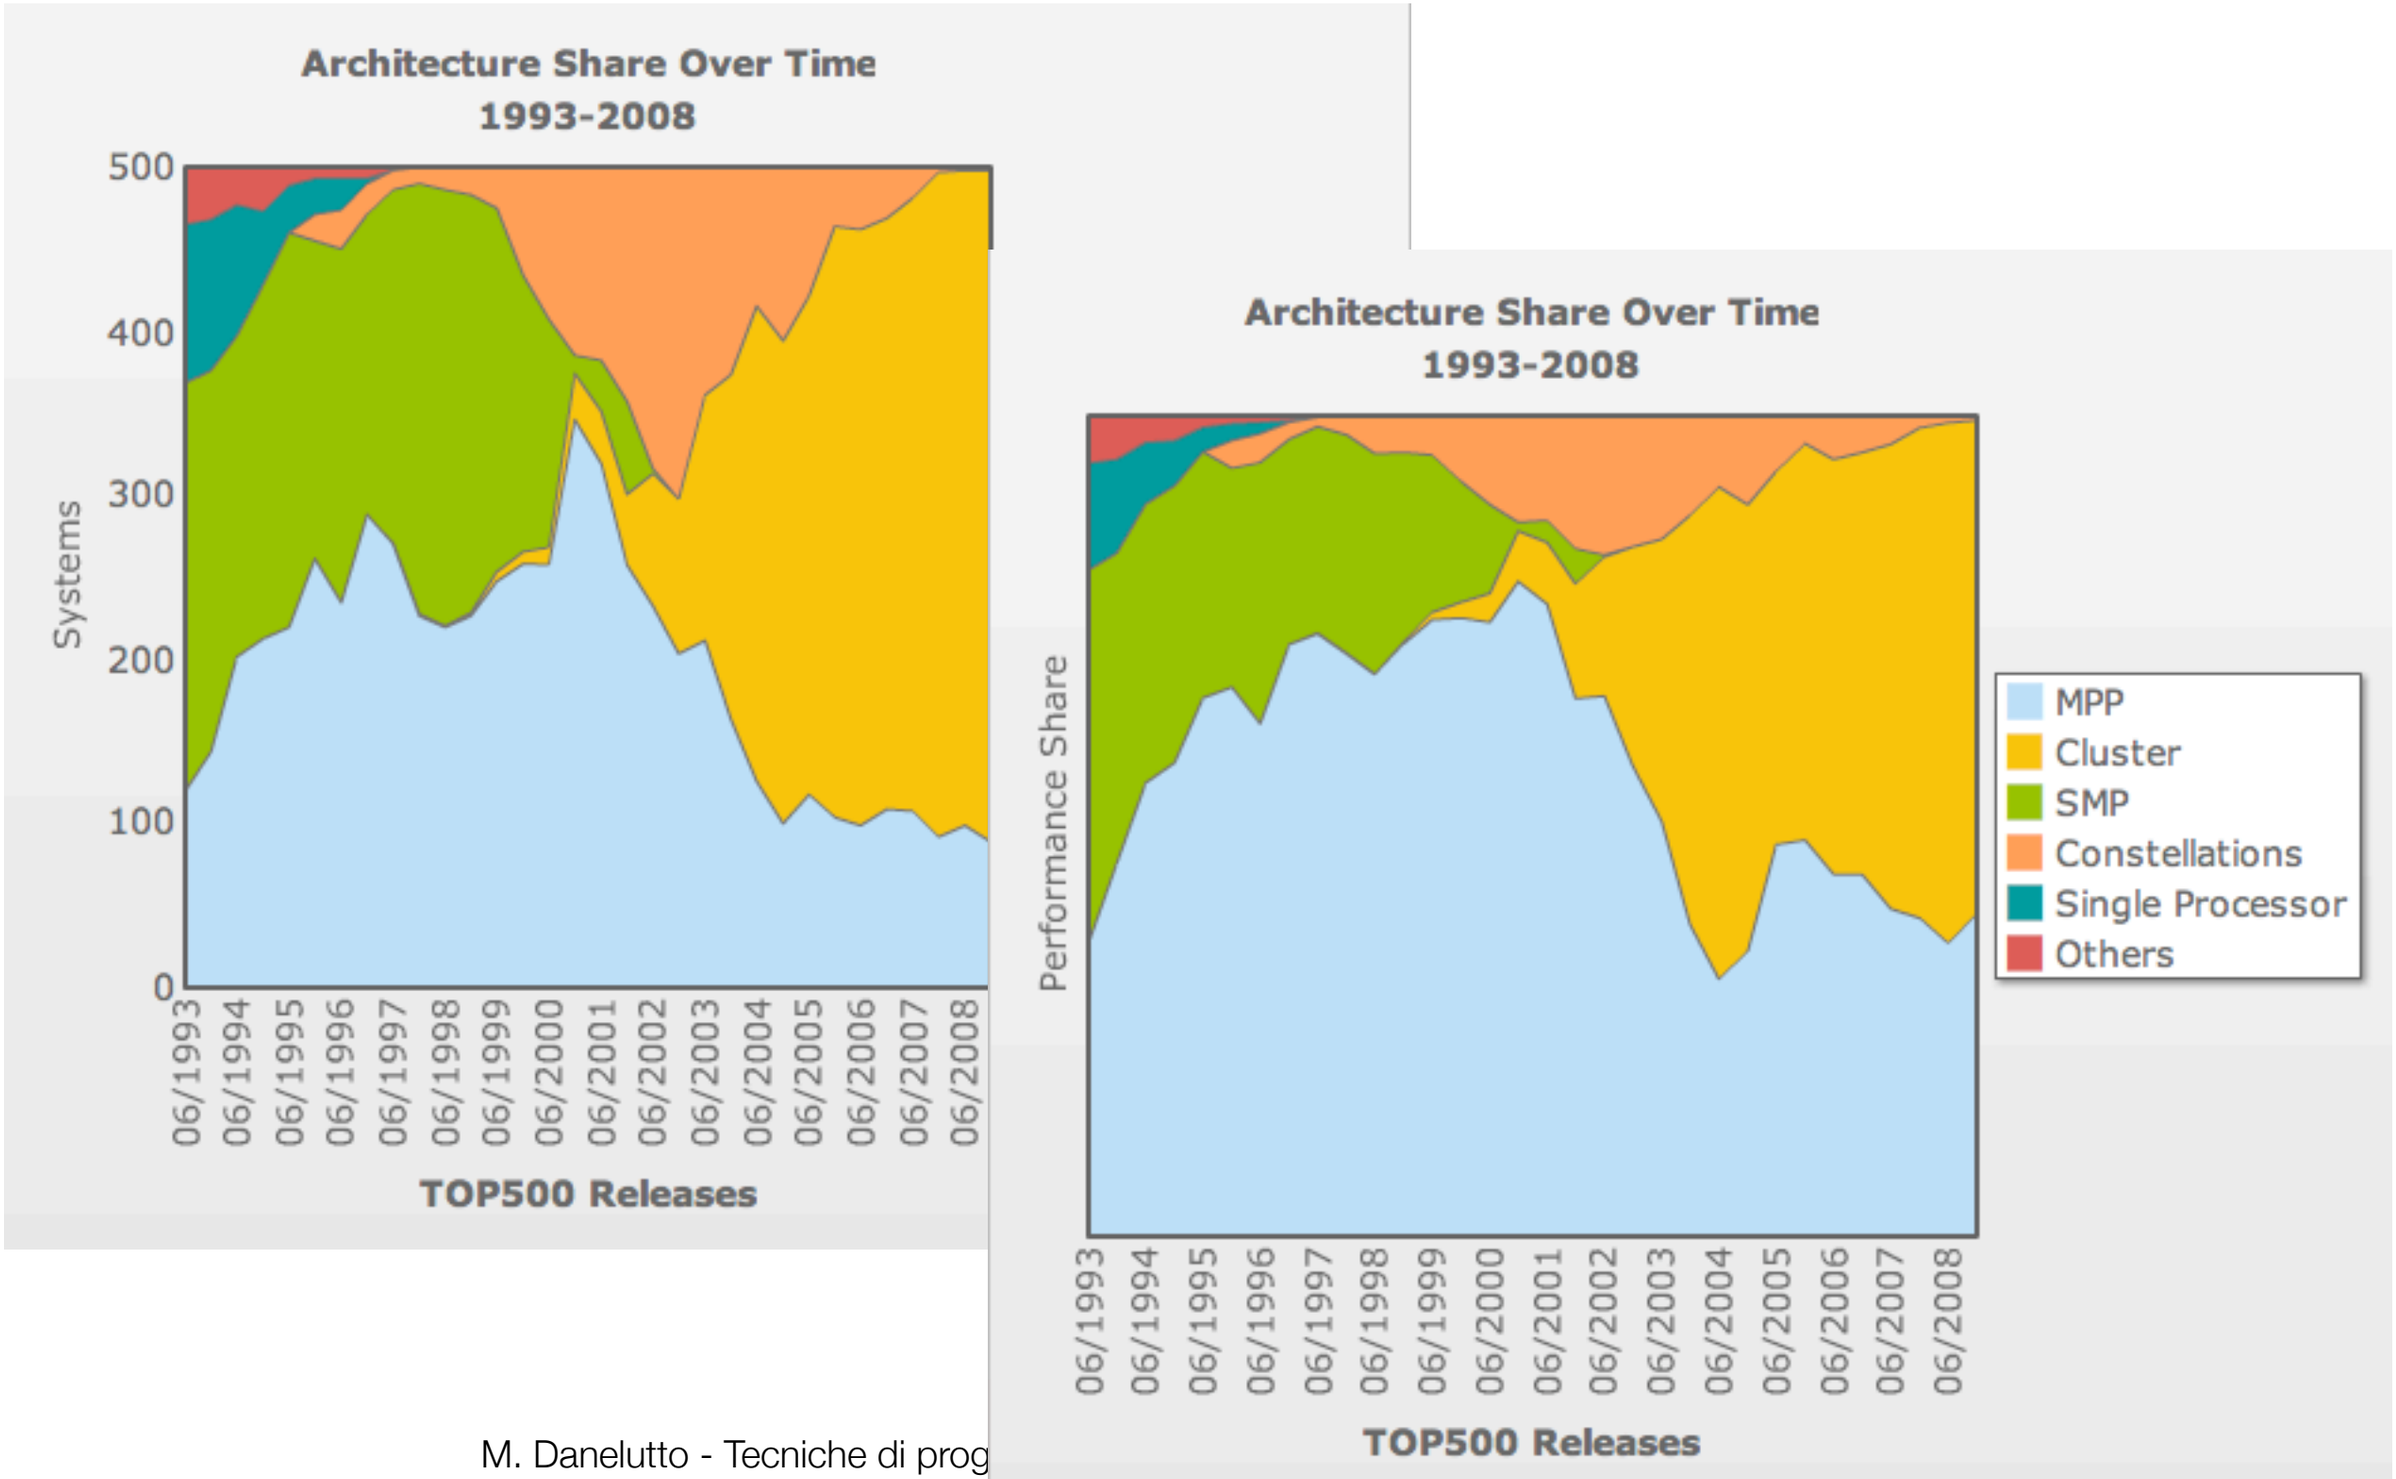

Evolution: processor architecture (top500)

- 1993

Processor Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Vector	334	66.80 %	650	792	1242
Scalar	131	26.20 %	408	859	15606
SIMD	35	7.00 %	64	135	54272
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

- 2000

Processor Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Vector	60	12.00 %	11685	13322	2648
Scalar	440	88.00 %	52545	82354	113010
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

- 2006

Processor Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Vector	7	1.40 %	85074	97805	8426
Scalar	493	98.60 %	3442384	5115743	1012535
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

Evolution: processor architecture (top500)

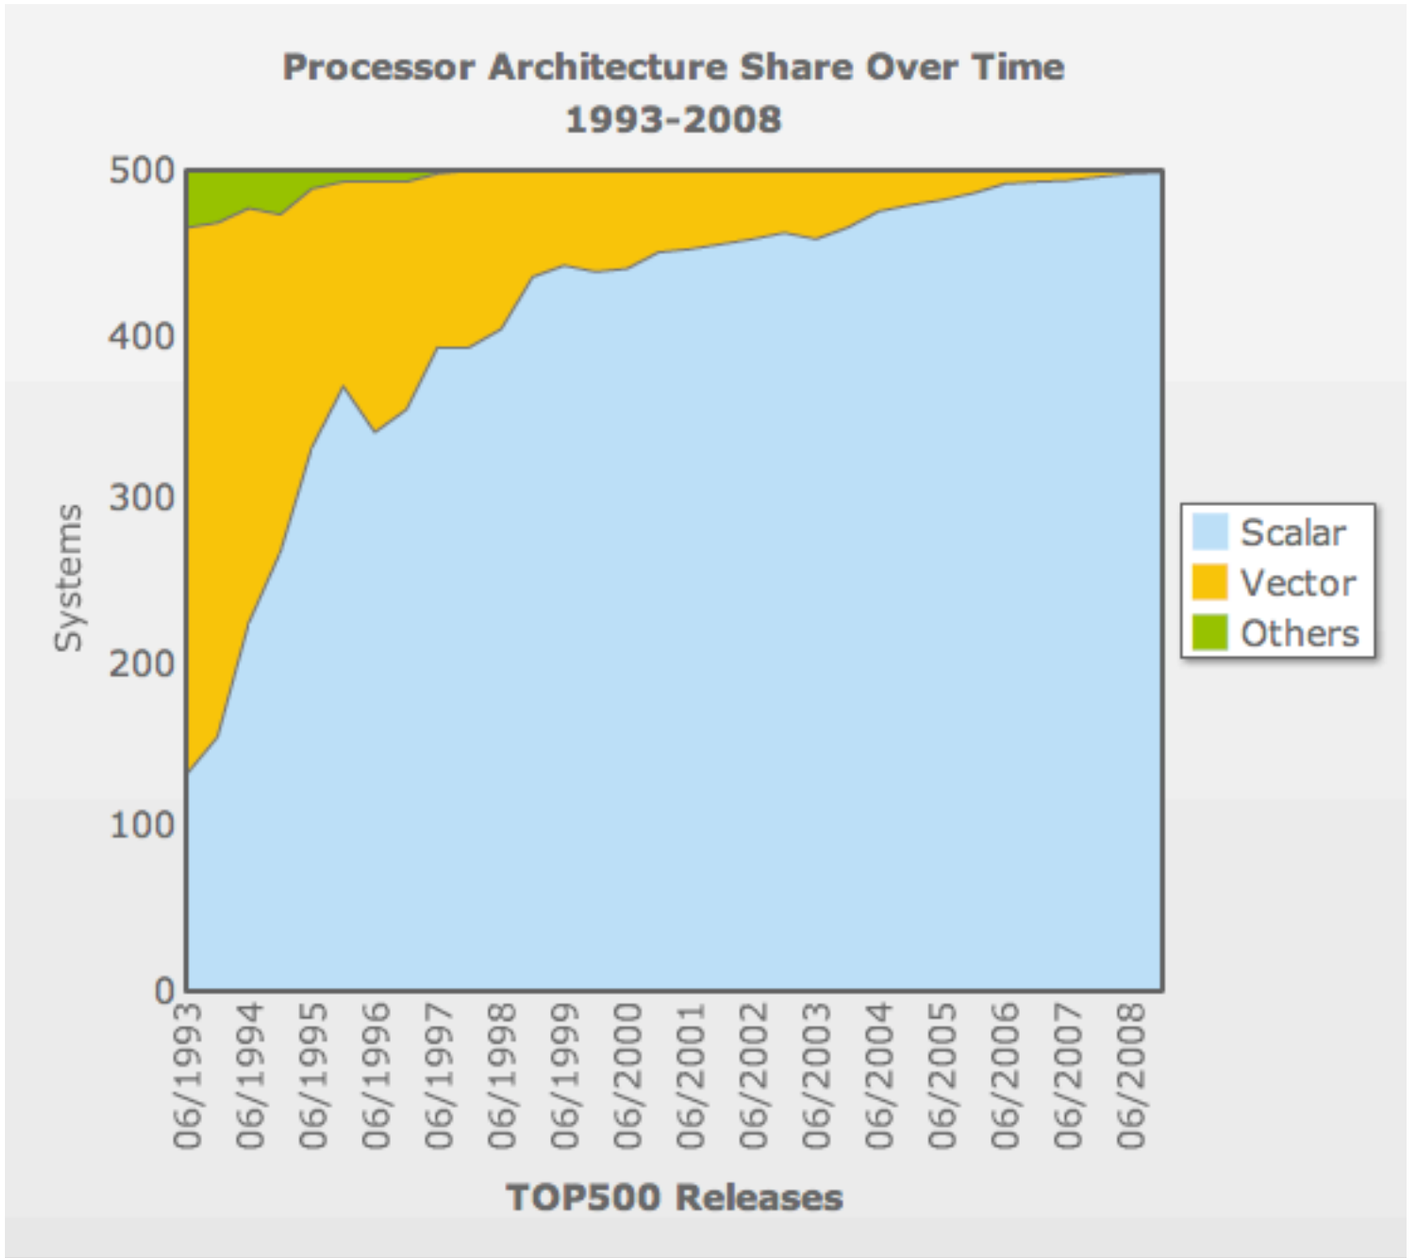

Evolution: processor architecture (top500)

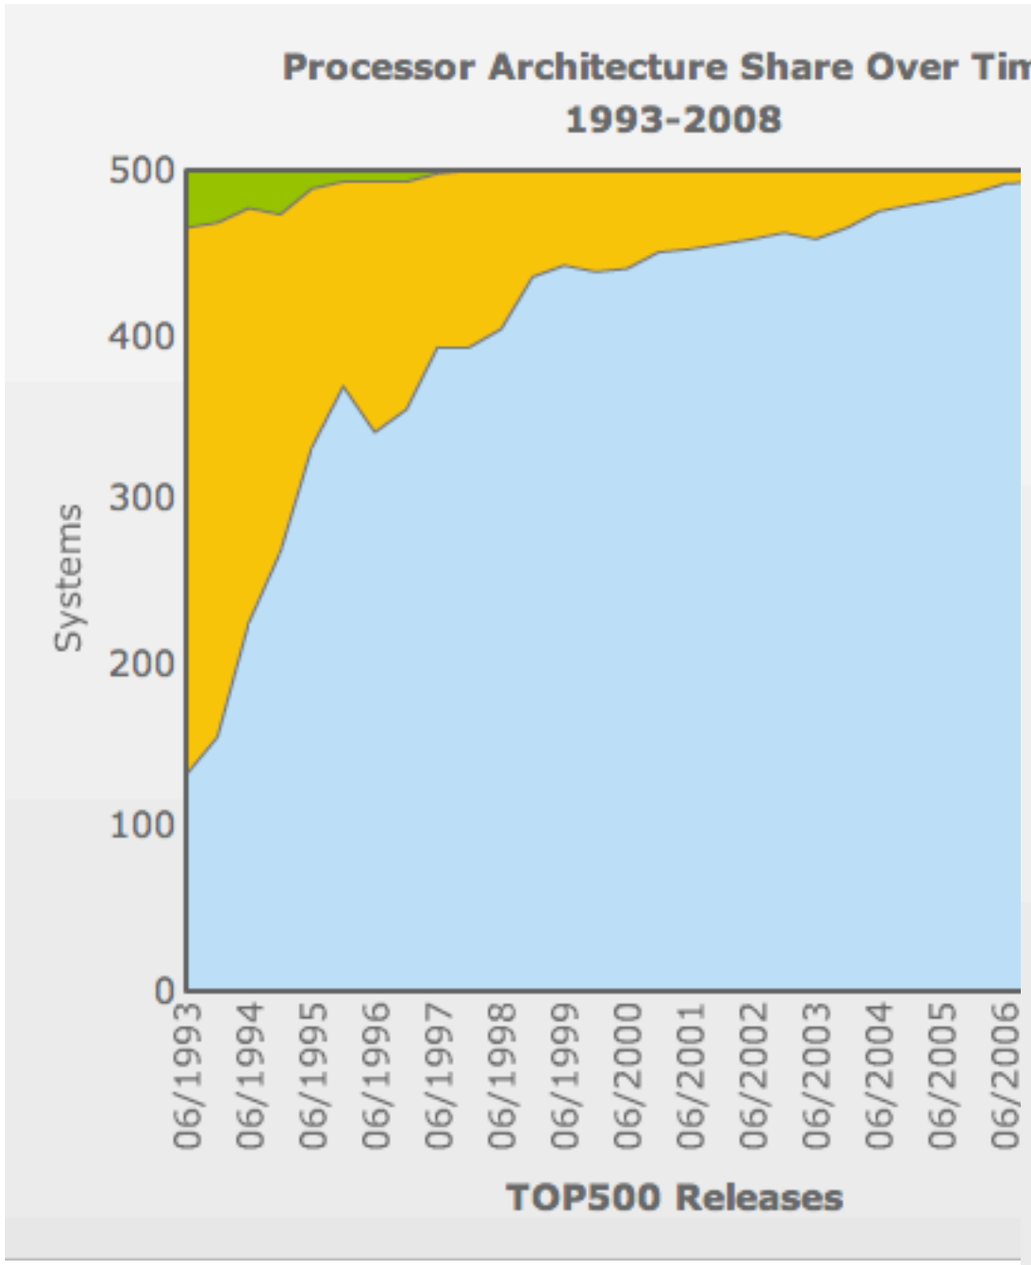

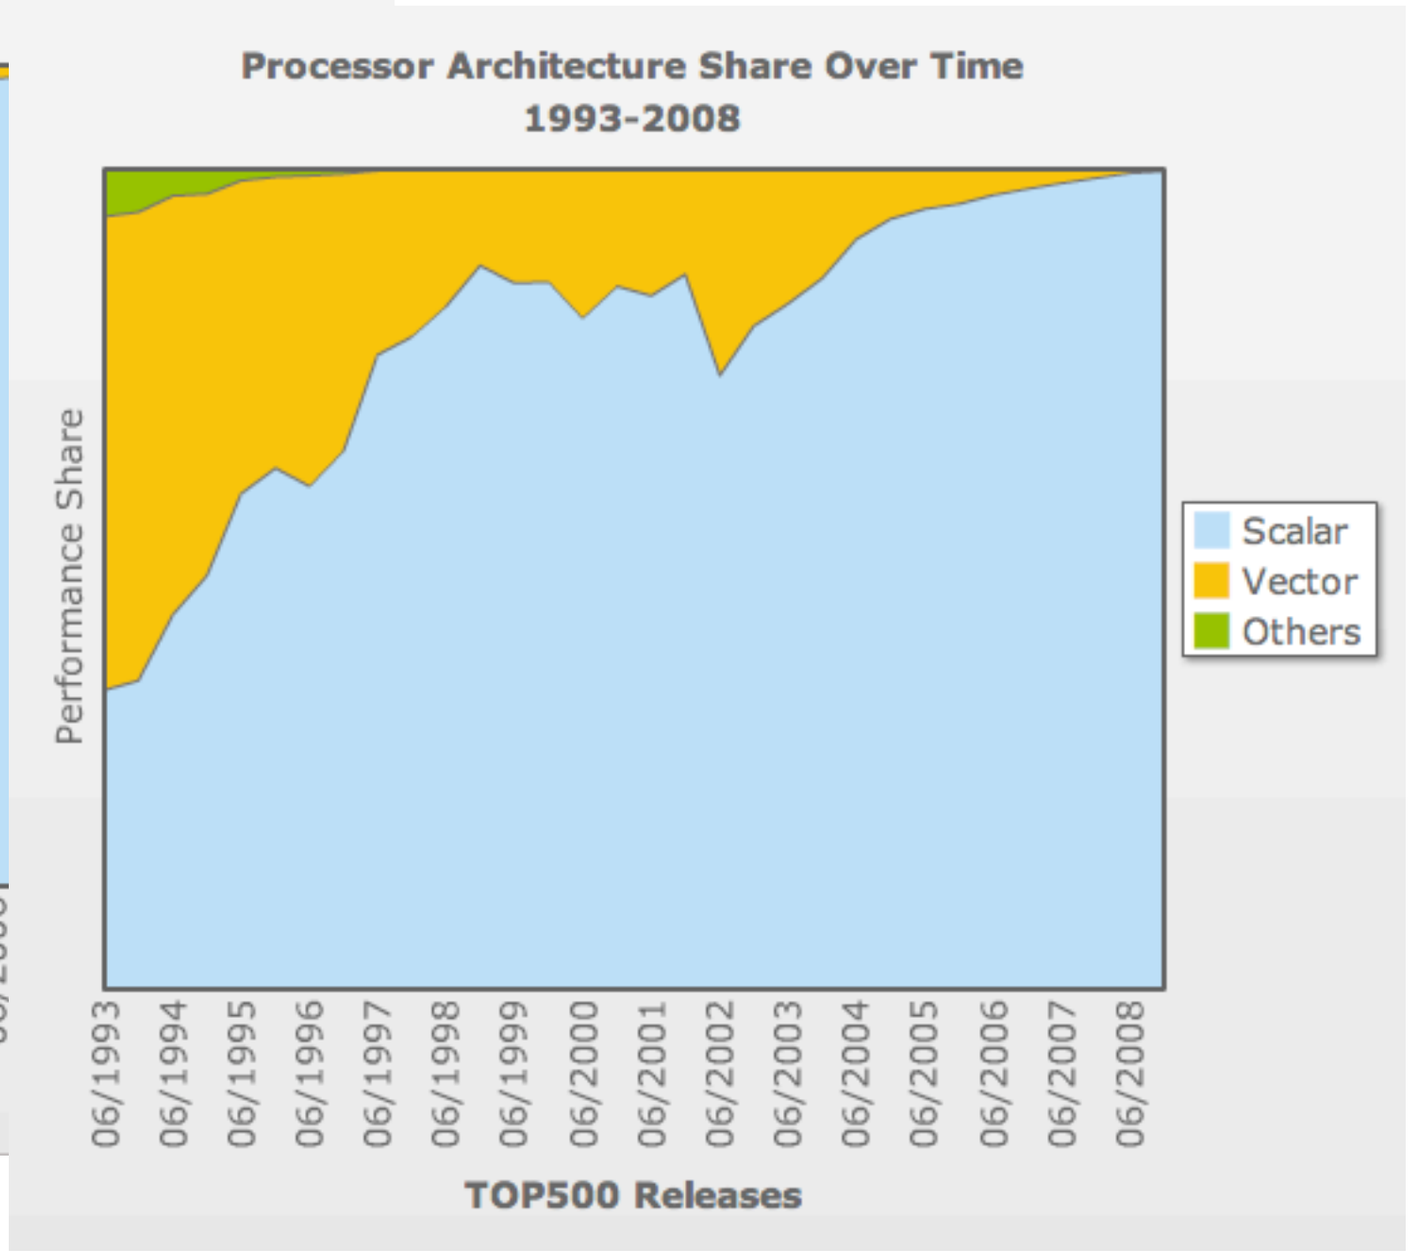

Evolution: interconnection family (top500)

- 1993

Interconnect Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
N/A	398	79.60 %	709	1007	22478
Crossbar	63	12.60 %	151	206	44162
Fat Tree	39	7.80 %	263	573	4480
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

- 2000

Interconnect Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
N/A	133	26.60 %	6851	8588	13398
Myrinet	42	8.40 %	2501	5905	3768
Quadrics	5	1.00 %	1535	2101	1576
Ethernet	6	1.20 %	309	726	1376
Fast Ethernet	3	0.60 %	159	418	652
Crossbar	115	23.00 %	16600	19957	17900
SP Switch	141	28.20 %	18327	31268	38728
HIPPI	1	0.20 %	1608	3072	6144
Proprietary	6	1.20 %	2635	3770	9952
Cray Interconnect	48	9.60 %	13706	19869	22164
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

- 2006

Interconnect Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Myrinet	79	15.80 %	420552	642430	110270
Quadrics	14	2.80 %	149888	192668	40420
Gigabit Ethernet	213	42.60 %	765828	1477104	243420
Infiniband	78	15.60 %	592273	841745	122136
Crossbar	11	2.20 %	112242	147150	16426
Mixed	5	1.00 %	74198	106008	16320
NUMAlink	17	3.40 %	130621	142743	22400
SP Switch	42	8.40 %	361555	495078	74916
Proprietary	30	6.00 %	669508	855210	311776
Cray Interconnect	9	1.80 %	244757	305978	61188
RapidArray	2	0.40 %	6037	7435	1689
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

Evolution: interconnection family (top500)

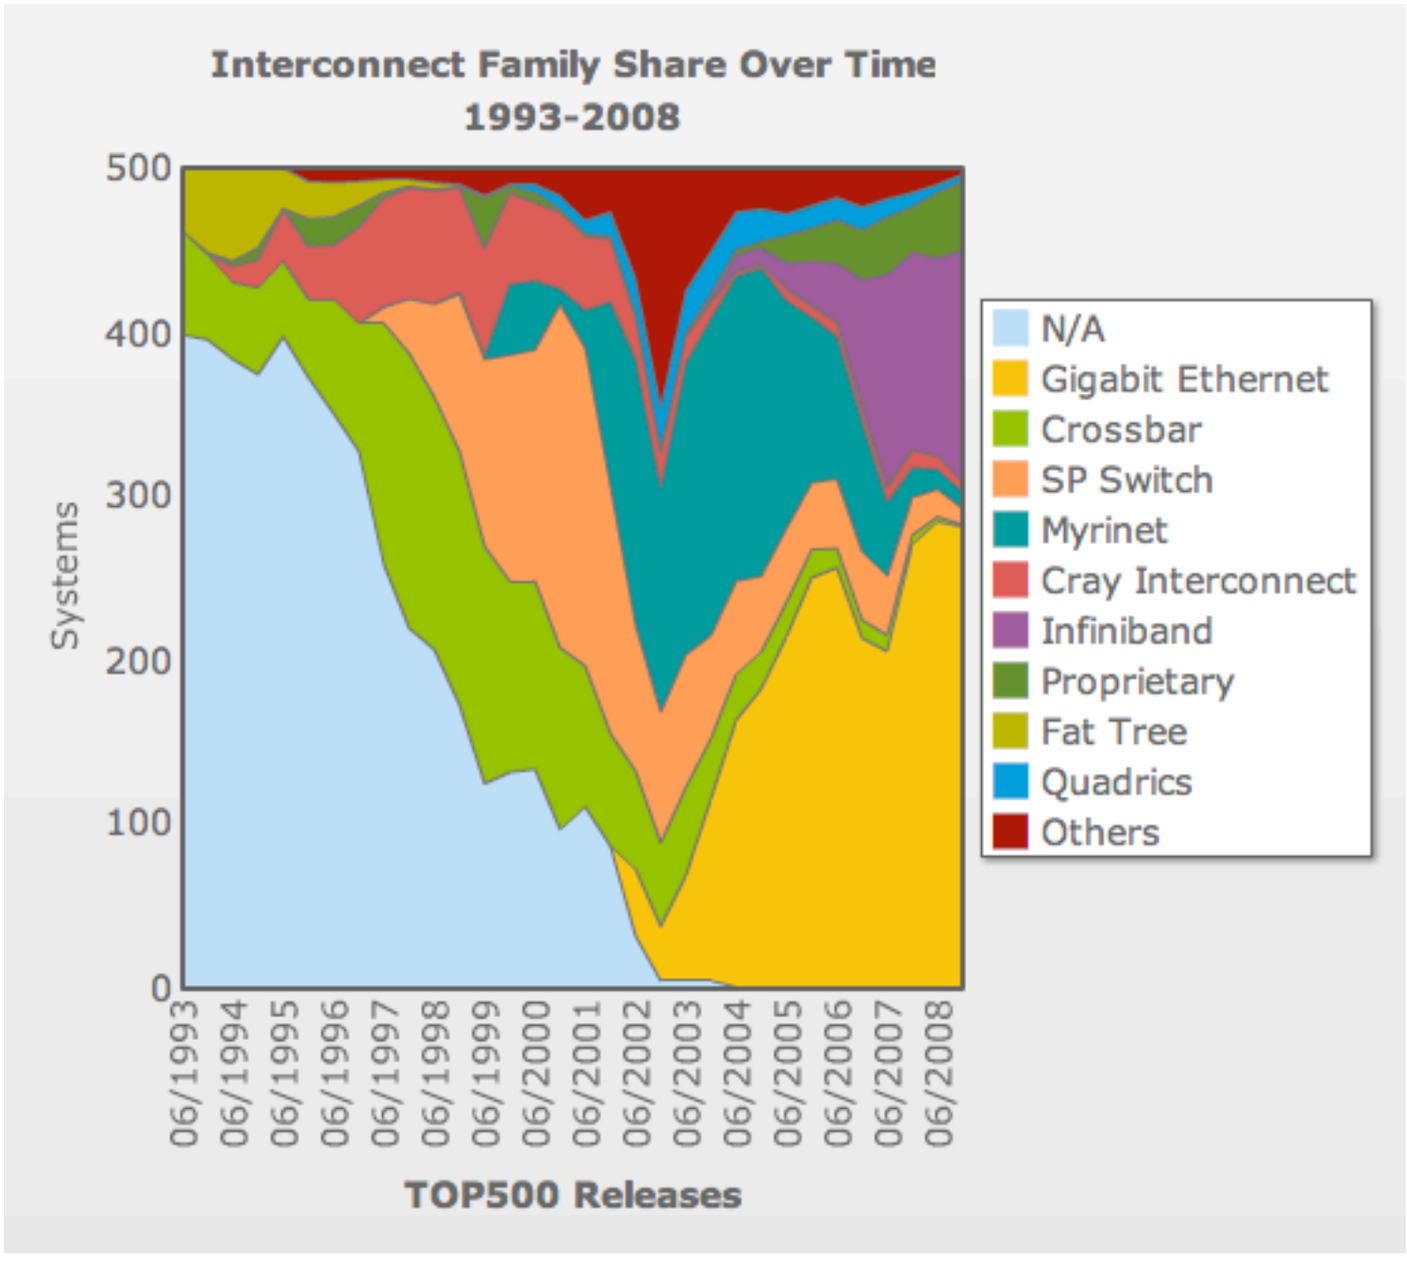

Evolution: interconnection family (top500)

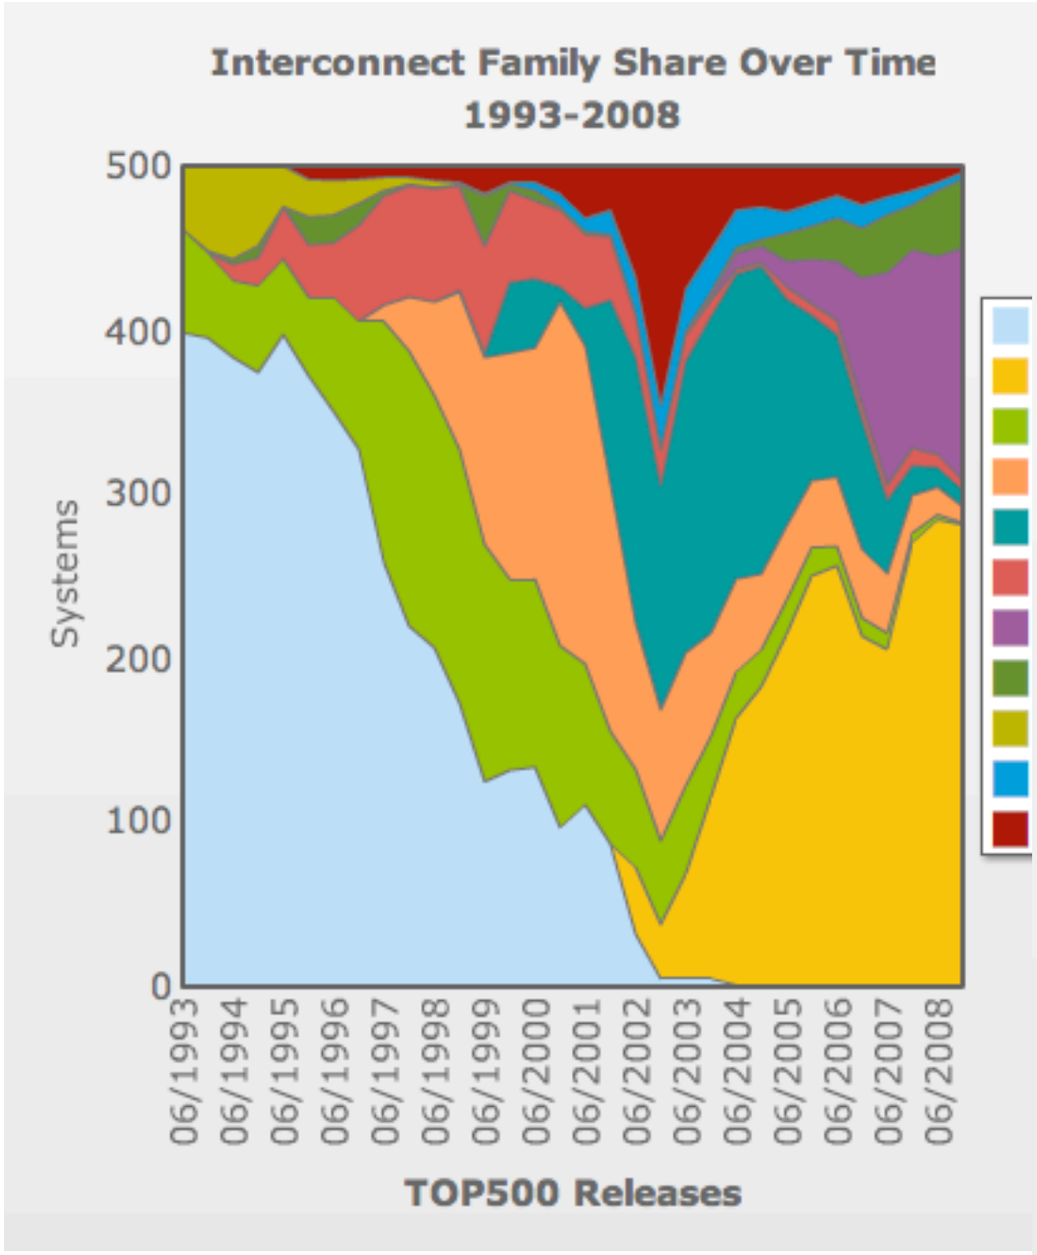

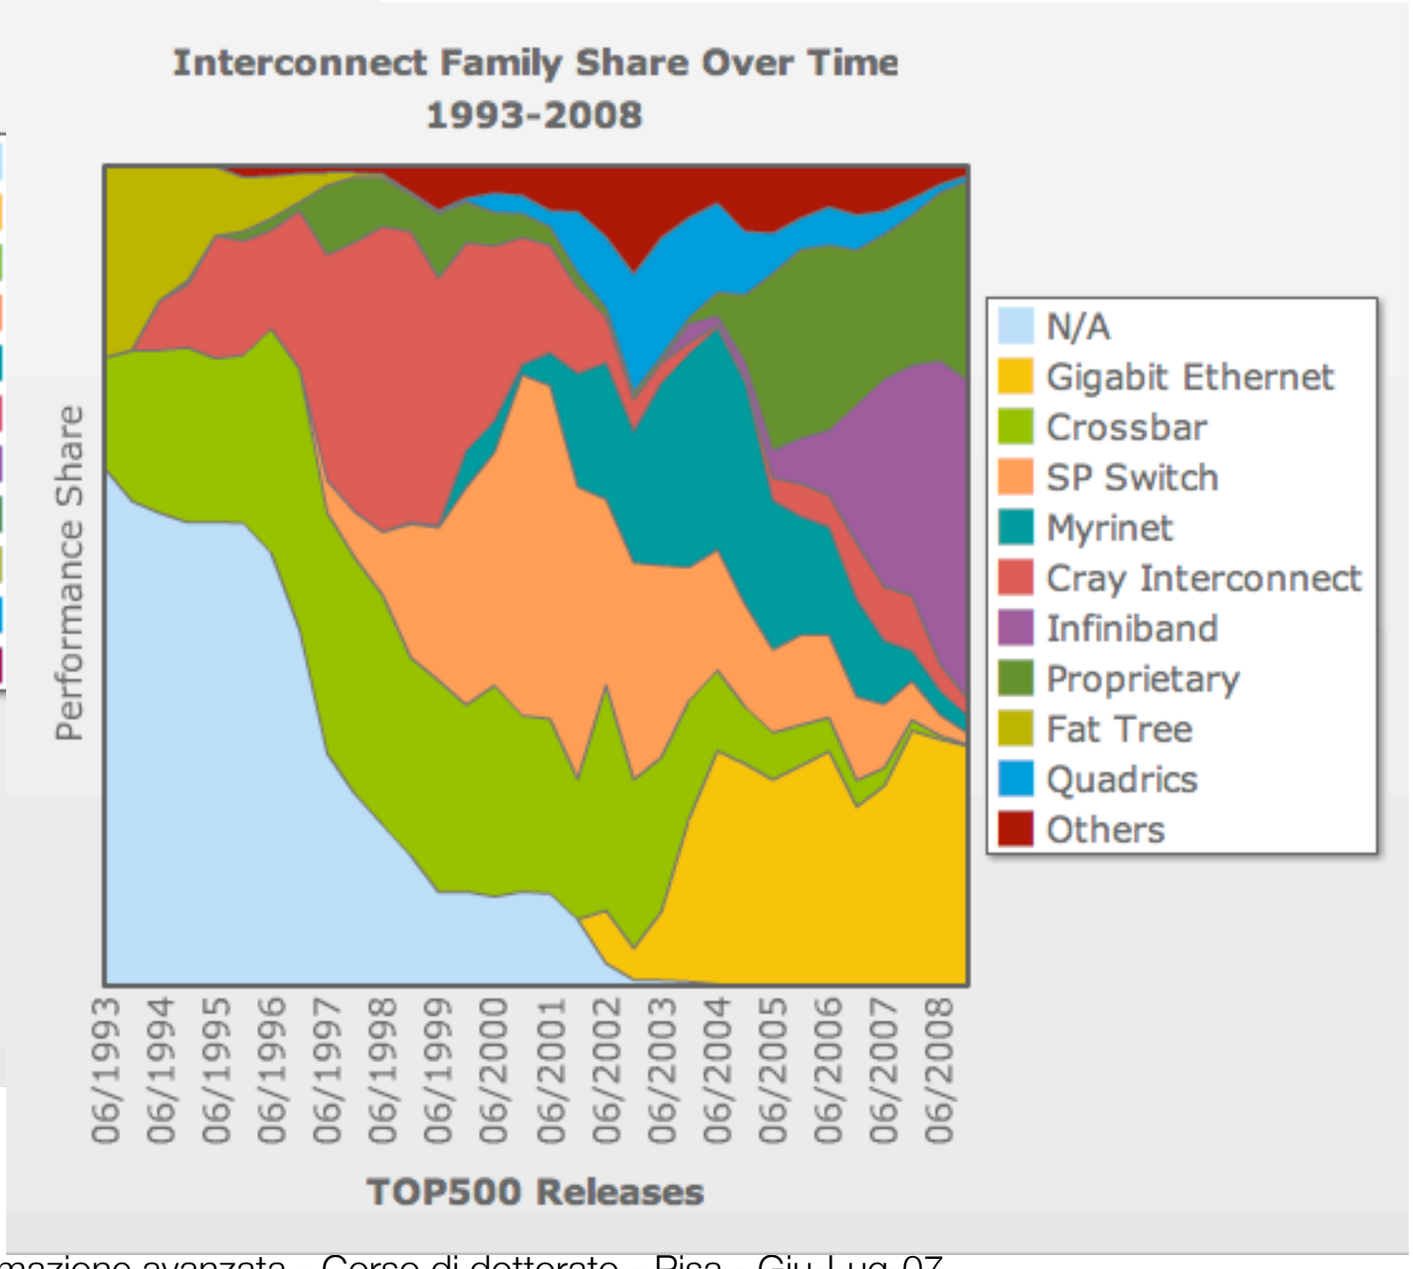

Evolution: operating system (top500)

- 1993

Operating system Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Unix	468	93.60 %	992	1640	71073
BSD Based	23	4.60 %	123	135	38
N/A	9	1.80 %	8	11	9
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

- 2000

Operating system Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Linux	28	5.60 %	1935	2763	4560
Unix	453	90.60 %	59805	89873	110064
BSD Based	17	3.40 %	2197	2320	314
N/A	2	0.40 %	294	720	720
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

- 2006

Operating system Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Linux	376	75.20 %	2014910	3195766	516189
Unix	86	17.20 %	559636	807423	142104
BSD Based	3	0.60 %	47697	53248	5888
Mixed	32	6.40 %	872226	1104103	350484
Mac OS	3	0.60 %	32989	53008	6296
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

Evolution: operating system (top500)

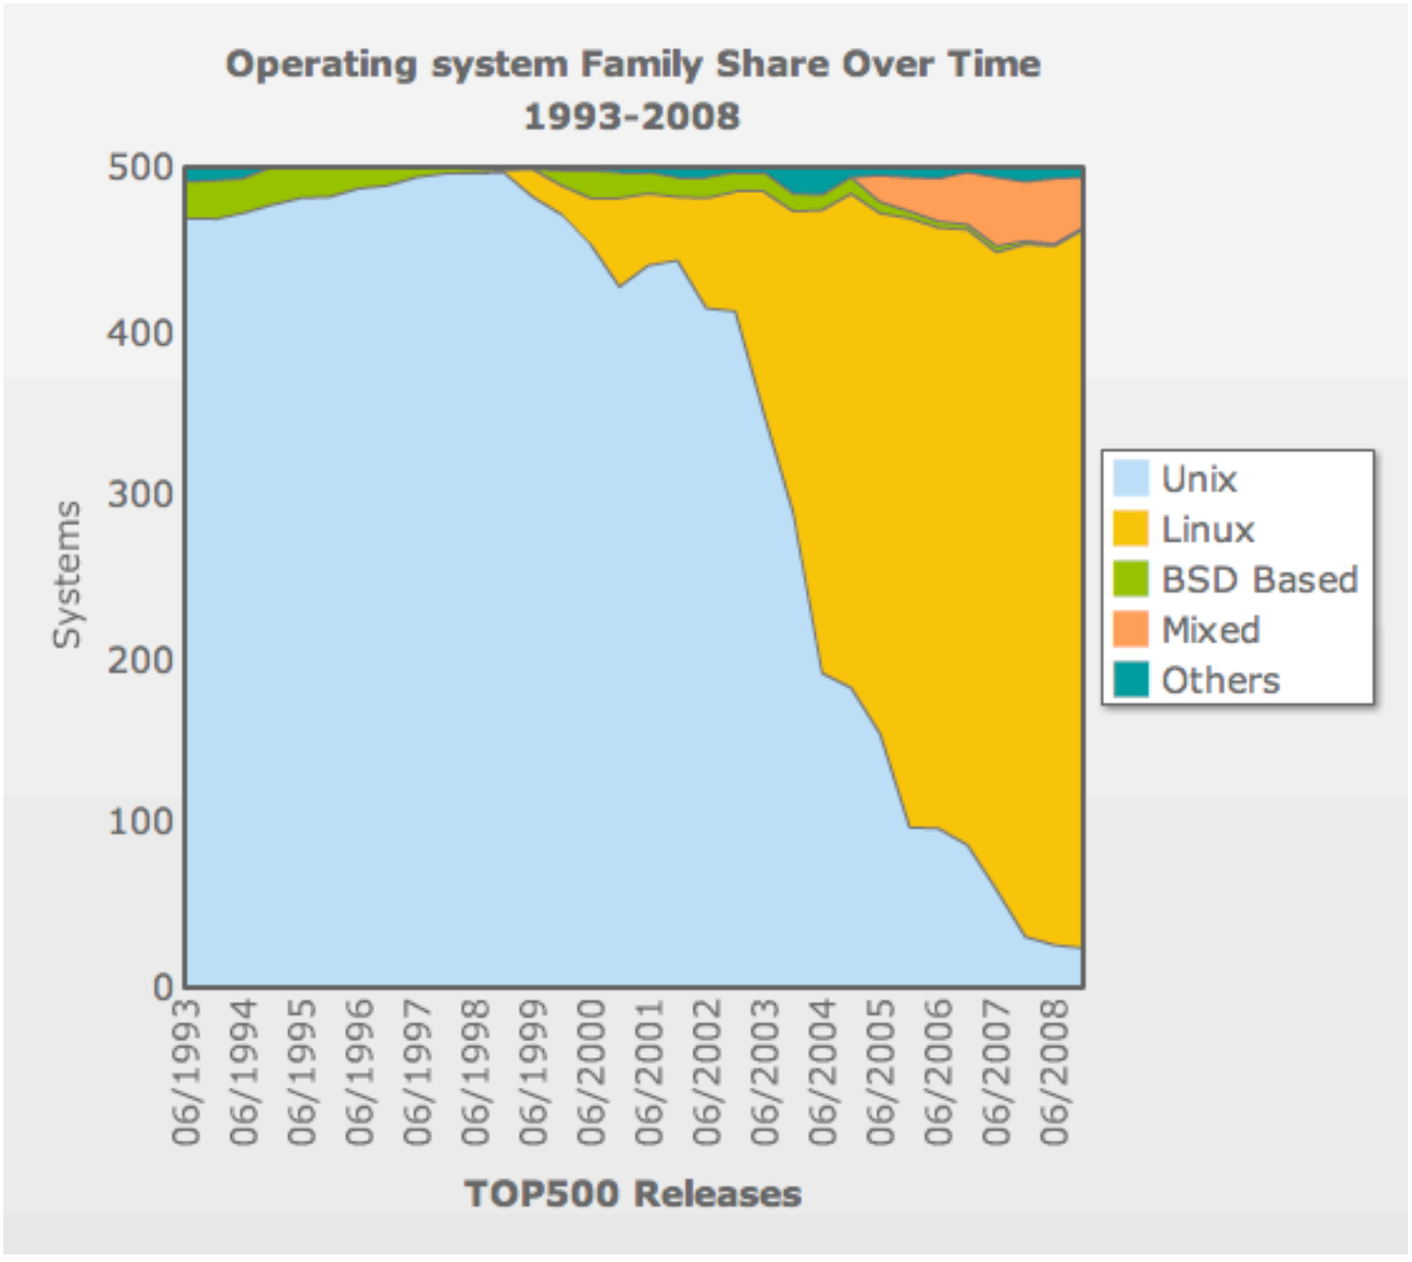

Evolution: operating system (top500)

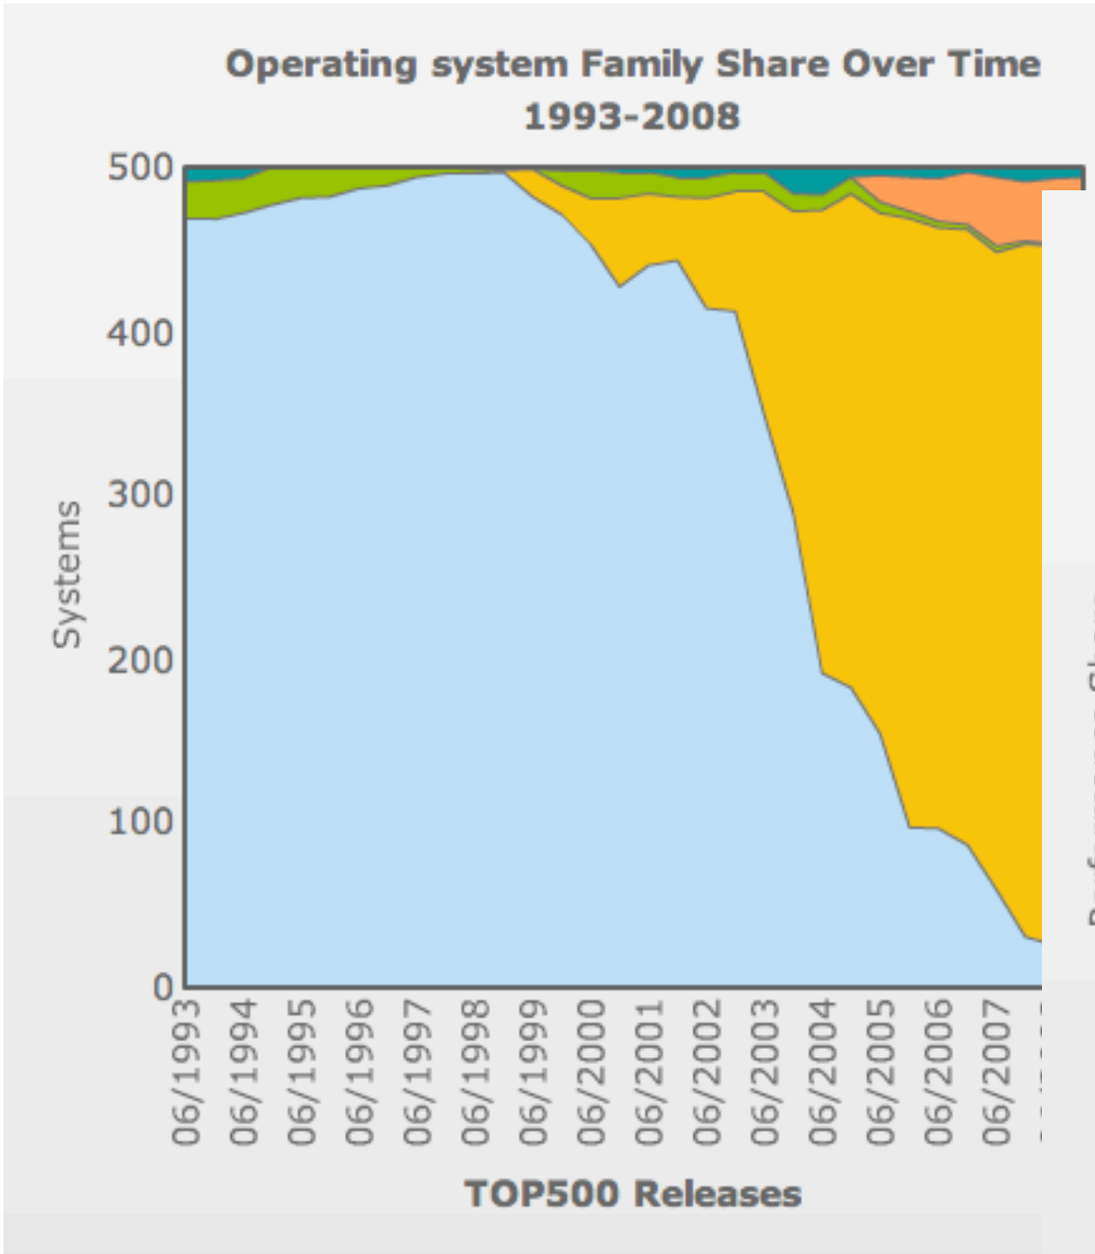

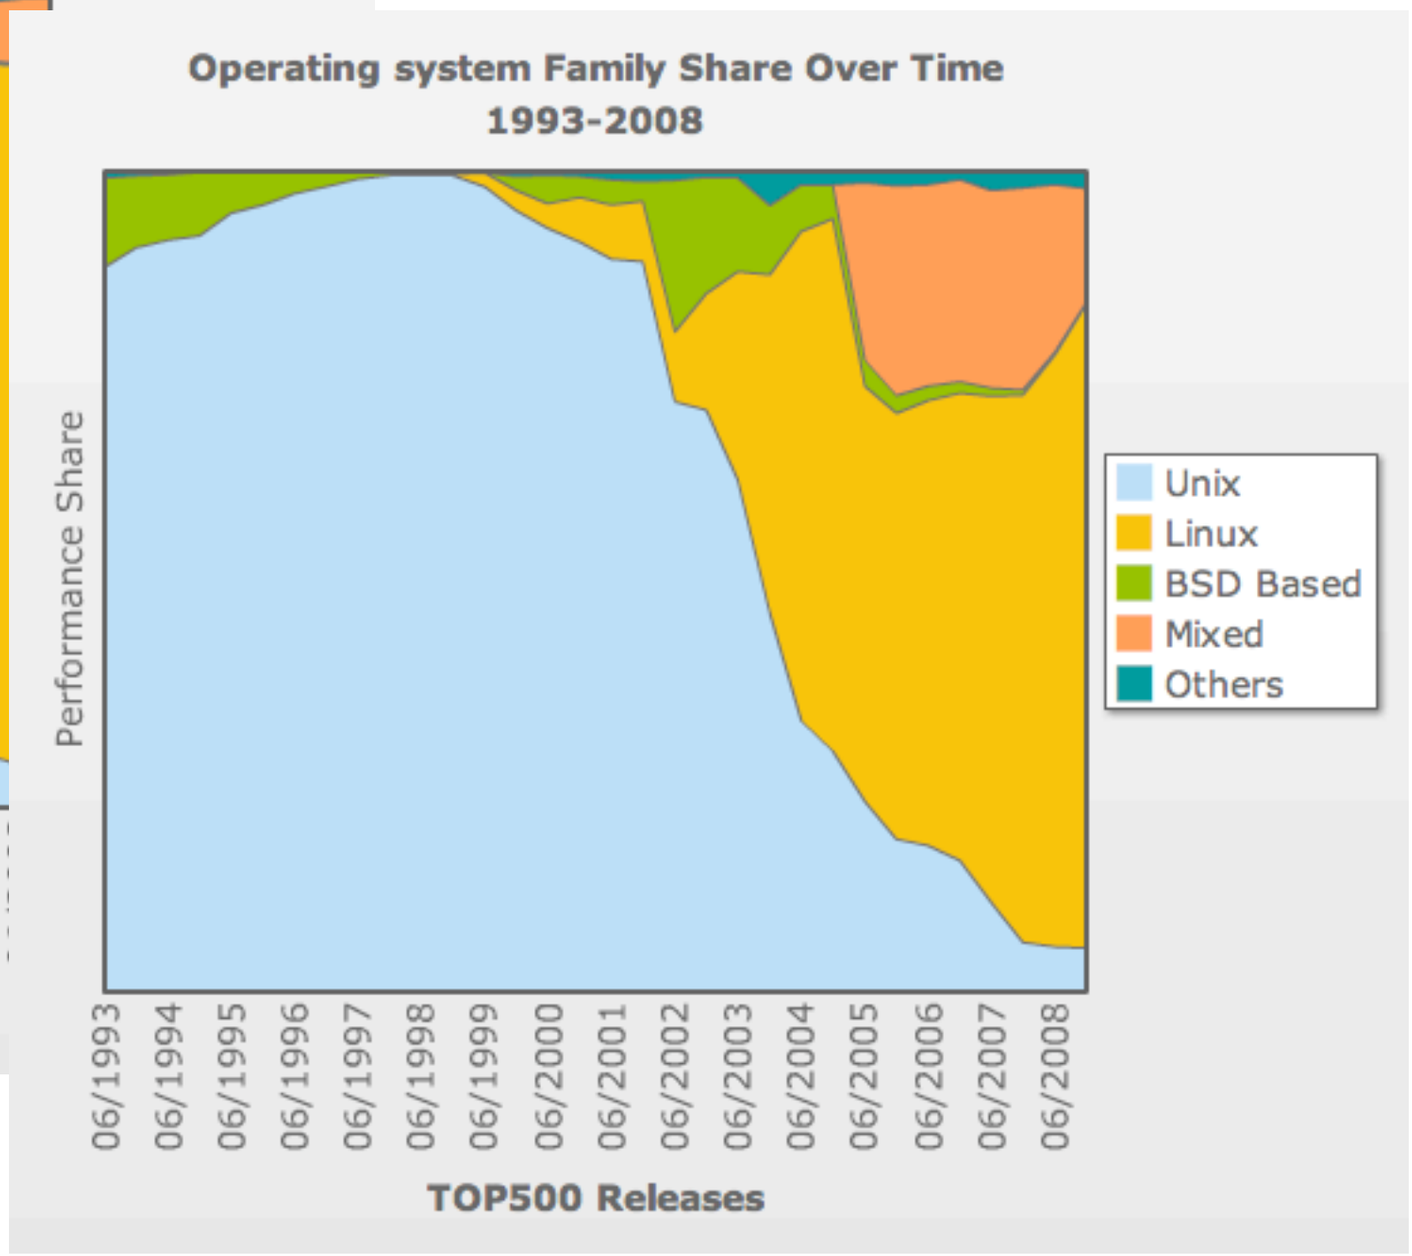

Evolution: operating system (top500)

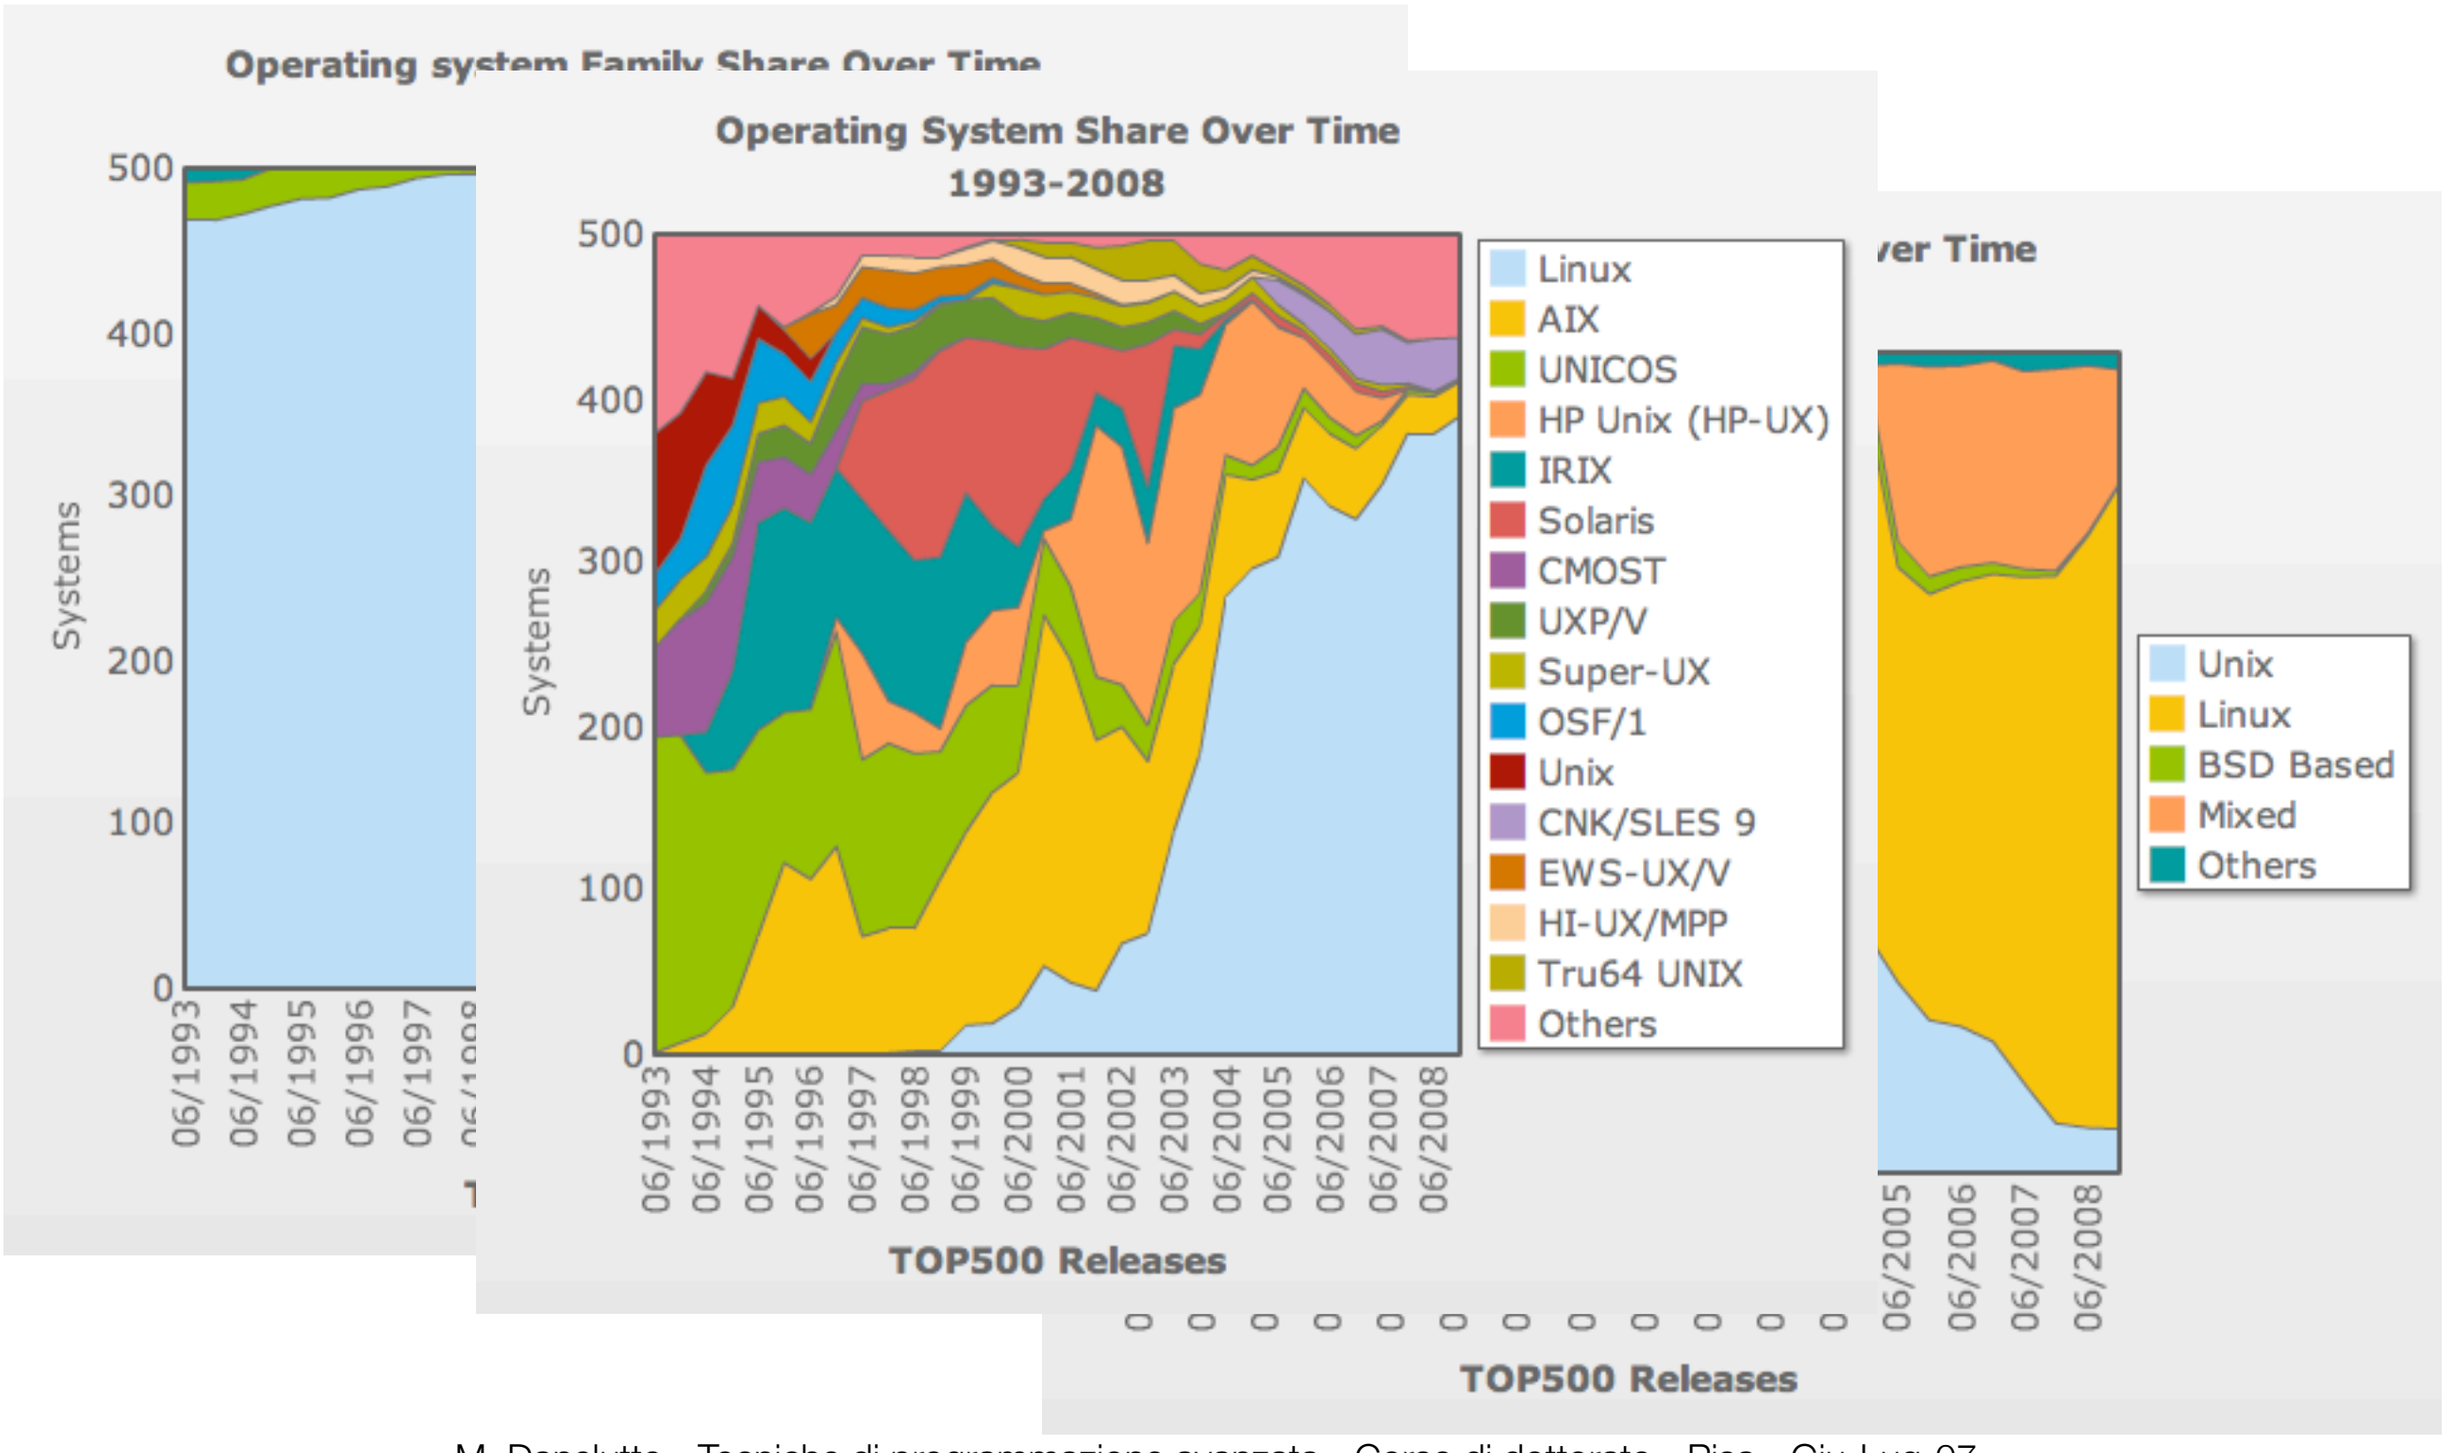

Evolution: number of processors (top500)

- 1993
- 2000
- 2006

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
1	95	19.00 %	140	176	95
2	72	14.40 %	83	94	144
3-4	98	19.60 %	144	194	380
5-8	57	11.40 %	119	145	431
9-16	25	5.00 %	172	197	400
17-32	33	6.60 %	47	96	1052
33-64	38	7.60 %	70	136	2274
65-128	30	6.00 %	81	172	3036
129-256	8	1.60 %	33	85	1836
257-512	13	2.60 %	90	184	6656
513-1024	8	1.60 %	102	224	7712
1025-2048	23	4.60 %	42	84	47104
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

Evolution: number of processors (top500)

- 1993

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
1	95	19.00 %	140	176	95
2	72	14.40 %	83	94	144
3-4	98	19.60 %	144	194	380

- 2000

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
5-8	6	1.20 %	299	320	40
9-16	10	2.00 %	1195	1248	158
17-32	19	3.80 %	2270	2452	550
33-64	161	32.20 %	10680	14405	9887
65-128	119	23.80 %	10392	13880	13356
129-256	104	20.80 %	8508	14660	20141
257-512	43	8.60 %	7899	12753	15632
513-1024	23	4.60 %	8927	13229	15948
1025-2048	11	2.20 %	7825	12468	15860
2049-4096	1	0.20 %	104	125	2502
4k-8k	2	0.40 %	3752	6928	11952
8k-16k	1	0.20 %	2379	3207	9632
Totals	500	100%	64230.11	95676.11	115658

- 2006

Evolution: number of processors (top500)

- 1993

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
1	95	19.00 %	140	176	95
2	72	14.40 %	83	94	144
3-4	98	19.60 %	144	194	380

- 2000

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
5-8	6	1.20 %	299	320	40
9-16	10	2.00 %	1195	1248	158
17-32	19	3.80 %	2270	2452	550
33-64	161	32.20 %	10680	14405	9887
65-128	119	23.80 %	10392	13880	13356
129-256	104	20.80 %	8508	14660	20141
257-512	43	8.60 %	7899	12753	15632
513-1024	23	4.60 %	8927	13229	15948
1025-2048	11	2.20 %	7825	12468	15860
2049-4096	1	0.20 %	104	125	2502
4k-8k	2	0.40 %	3752	6928	11952
8k-16k	1	0.20 %	2379	3207	9632
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

- 2006

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
33-64	4	0.80 %	16204	22989	196
65-128	2	0.40 %	18072	21504	160
129-256	3	0.60 %	9630	10854	696
257-512	36	7.20 %	117710	162157	17836
513-1024	192	38.40 %	716626	1144890	171117
1025-2048	185	37.00 %	865405	1423050	262844
2049-4096	38	7.60 %	372432	584622	98884
4k-8k	19	3.80 %	357877	470662	95140
8k-16k	17	3.40 %	542883	717847	159128
16k-32k	2	0.40 %	138730	173286	42928
32k-64k	1	0.20 %	91290	114688	40960
64k-128k	1	0.20 %	280600	367000	131072
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

Evolution: number of processors (top500)

- 1993
- 2000
- 2006

Number of Processors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
1	95	19.00 %	140	176	95
2	72	14.40 %	83	94	144
3-4	98	19.60 %	144	194	380

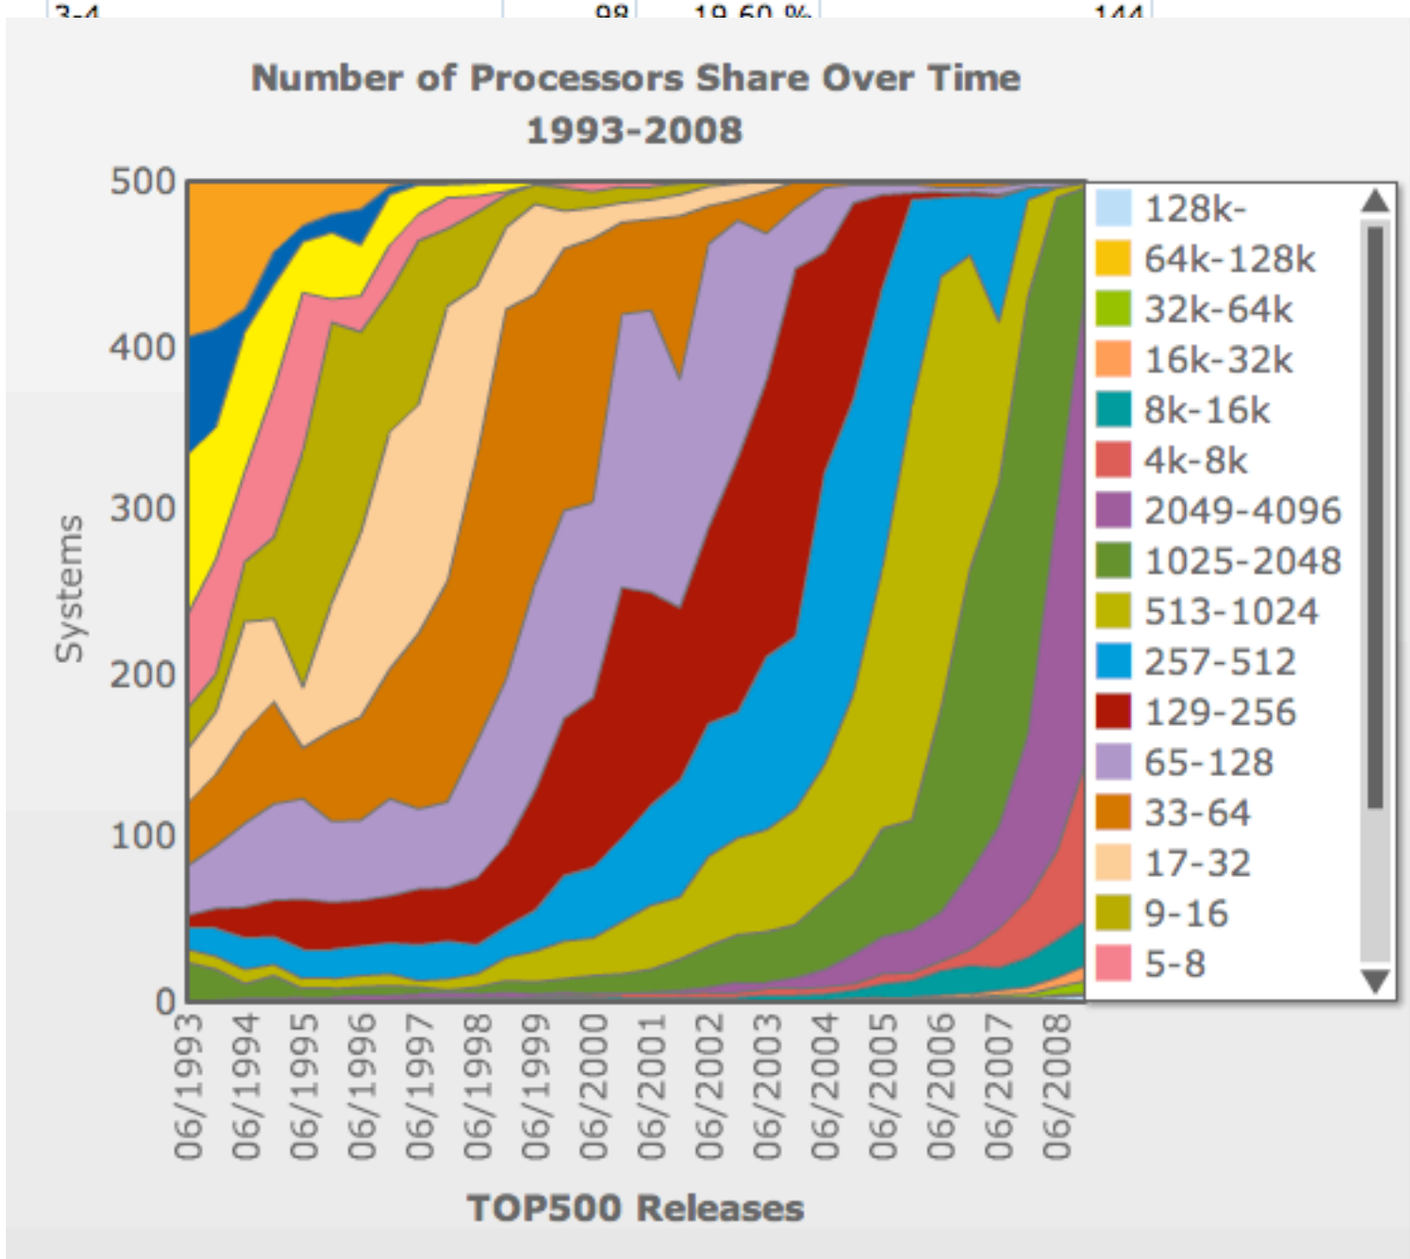

32k-64k	1	0.20 %	91290	114688	40960
64k-128k	1	0.20 %	280600	367000	131072
Totals	**500**	**100%**	**3527458.35**	**5213548.18**	**1020961**

F)	Processor Sum
320	40
1248	158
2452	550
.4405	9887
.3880	13356
.4660	20141
.2753	15632
.3229	15948
.2468	15860
125	2502
6928	11952
3207	9632
	115658

)	Processor Sum
989	196
504	160
854	696
157	17836
890	171117
050	262844
622	98884
662	95140
847	159128
286	42928

Evolution: country (top500)

- 1993

- 2000

- 2006

Evolution: country (top500)

Countries	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Australia	9	1.80 %	11	22	670
Austria	2	0.40 %	2	3	57
Brazil	2	0.40 %	4	4	3
Canada	3	0.60 %	23	25	7
Denmark	3	0.60 %	2	2	28
Finland	2	0.40 %	1	1	8
France	26	5.20 %	35	58	3019
Germany	59	11.80 %	69	102	8523
Greece	1	0.20 %	0	1	16
Hong Kong	1	0.20 %	0	1	16
Italy	6	1.20 %	4	7	326
Japan	111	22.20 %	208	296	12296
Korea, South	3	0.60 %	3	4	10
Mexico	1	0.20 %	1	1	4
Netherlands	6	1.20 %	9	11	566
Norway	3	0.60 %	4	7	2108
Slovenia	1	0.20 %	1	1	6
Spain	2	0.40 %	2	3	3
Sweden	2	0.40 %	1	1	8
Switzerland	4	0.80 %	8	9	12
Taiwan	3	0.60 %	3	3	14
United Kingdom	25	5.00 %	46	62	1049
United States	225	45.00 %	686	1162	42371
Totals	**500**	**100%**	**1122.84**	**1786.21**	**71120**

Evolution: country (top500)

Countries	Count	Share %
Australia	9	1.80 %
Austria	2	0.40 %
Brazil	2	0.40 %
Canada	3	0.60 %
Denmark	3	0.60 %
Finland	2	0.40 %
France	26	5.20 %
Germany	59	11.80 %
Greece	1	0.20 %
Hong Kong	1	0.20 %
Italy	6	1.20 %
Japan	111	22.20 %
Korea, South	3	0.60 %
Mexico	1	0.20 %
Netherlands	6	1.20 %
Norway	3	0.60 %
Slovenia	1	0.20 %
Spain	2	0.40 %
Sweden	2	0.40 %
Switzerland	4	0.80 %
Taiwan	3	0.60 %
United Kingdom	25	5.00 %
United States	225	45.00 %
Totals	**500**	**100%**

Countries	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Australia	3	0.60 %	227	256	176
Brazil	1	0.20 %	79	102	128
Canada	9	1.80 %	729	941	856
China	2	0.40 %	100	135	160
Denmark	1	0.20 %	51	72	48
Finland	2	0.40 %	346	482	668
France	20	4.00 %	2166	3043	3063
Germany	65	13.00 %	7119	11184	11056
Greece	1	0.20 %	44	106	48
Italy	6	1.20 %	456	688	780
Japan	61	12.20 %	9804	11956	11067
Korea, South	4	0.80 %	455	527	282
Luxembourg	6	1.20 %	300	358	448
Mexico	4	0.80 %	221	275	344
Netherlands	5	1.00 %	360	620	890
New Zealand	1	0.20 %	109	158	132
Peru	1	0.20 %	45	51	64
Poland	1	0.20 %	59	77	96
Saudia Arabia	2	0.40 %	244	354	236
Singapore	1	0.20 %	46	48	6
Spain	2	0.40 %	89	102	128
Sweden	5	1.00 %	394	580	770
Switzerland	8	1.60 %	420	688	852
Taiwan	2	0.40 %	310	396	183
United Kingdom	28	5.60 %	3913	5498	5530
United States	259	51.80 %	36143	56978	77647
Totals	**500**	**100%**	**64230.11**	**95676.11**	**115658**

Evolution: country (top500)

Countries	Count	Share %
Australia	9	1.80 %
Austria	2	0.40 %
Brazil	2	0.40 %
Canada	3	0.60 %
Denmark	3	0.60 %
Finland	2	0.40 %
France	26	5.20 %
Germany	59	11.80 %
Greece	1	0.20 %
Hong Kong	1	0.20 %
Italy	6	1.20 %
Japan	111	22.20 %
Korea, South	3	0.60 %
Mexico	1	0.20 %
Netherlands	6	1.20 %
Norway	3	0.60 %
Slovenia	1	0.20 %
Spain	2	0.40 %
Sweden	2	0.40 %
Switzerland	4	0.80 %
Taiwan	3	0.60 %
United Kingdom	25	5.00 %
United States	225	45.00 %
Totals	500	100%

Countries	Count	Share %
Australia	3	0.60 %
Brazil	1	0.20 %
Canada	9	1.80 %
China	2	0.40 %
Denmark	1	0.20 %
Finland	2	0.40 %
France	20	4.00 %
Germany	65	13.00 %
Greece	1	0.20 %
Italy	6	1.20 %
Japan	61	12.20 %
Korea, South	4	0.80 %
Luxembourg	6	1.20 %
Mexico	4	0.80 %
Netherlands	5	1.00 %
New Zealand	1	0.20 %
Peru	1	0.20 %
Poland	1	0.20 %
Saudia Arabia	2	0.40 %
Singapore	1	0.20 %
Spain	2	0.40 %
Sweden	5	1.00 %
Switzerland	8	1.60 %
Taiwan	2	0.40 %
United Kingdom	28	5.60 %
United States	259	51.80 %
Totals	500	100%

Countries	Count	Share %	Rmax Sum (GF)
Australia	4	0.80 %	20671
Belgium	1	0.20 %	4493
Brazil	4	0.80 %	13668
Canada	8	1.60 %	37317
China	18	3.60 %	72192
Denmark	1	0.20 %	2791
Finland	1	0.20 %	8200
France	12	2.40 %	99870
Germany	18	3.60 %	145407
India	10	2.00 %	34162
Ireland	1	0.20 %	3142
Israel	2	0.40 %	7510
Italy	8	1.60 %	39172
Japan	30	6.00 %	286674
Korea, South	6	1.20 %	33715
Malaysia	3	0.60 %	12125
Mexico	1	0.20 %	5090
Netherlands	2	0.40 %	31114
New Zealand	1	0.20 %	3755
Norway	3	0.60 %	17473
Russia	2	0.40 %	9705
Saudia Arabia	4	0.80 %	11966
Singapore	2	0.40 %	6324
South Africa	2	0.40 %	5696
Spain	7	1.40 %	91600
Sweden	1	0.20 %	4999
Switzerland	5	1.00 %	47682
Taiwan	2	0.40 %	5535
Turkey	1	0.20 %	3288
United Arab Emirates	1	0.20 %	4713
United Kingdom	30	6.00 %	186420
United States	309	61.80 %	2270990
Totals	500	100%	3527458.35

Evolution: country (top500)

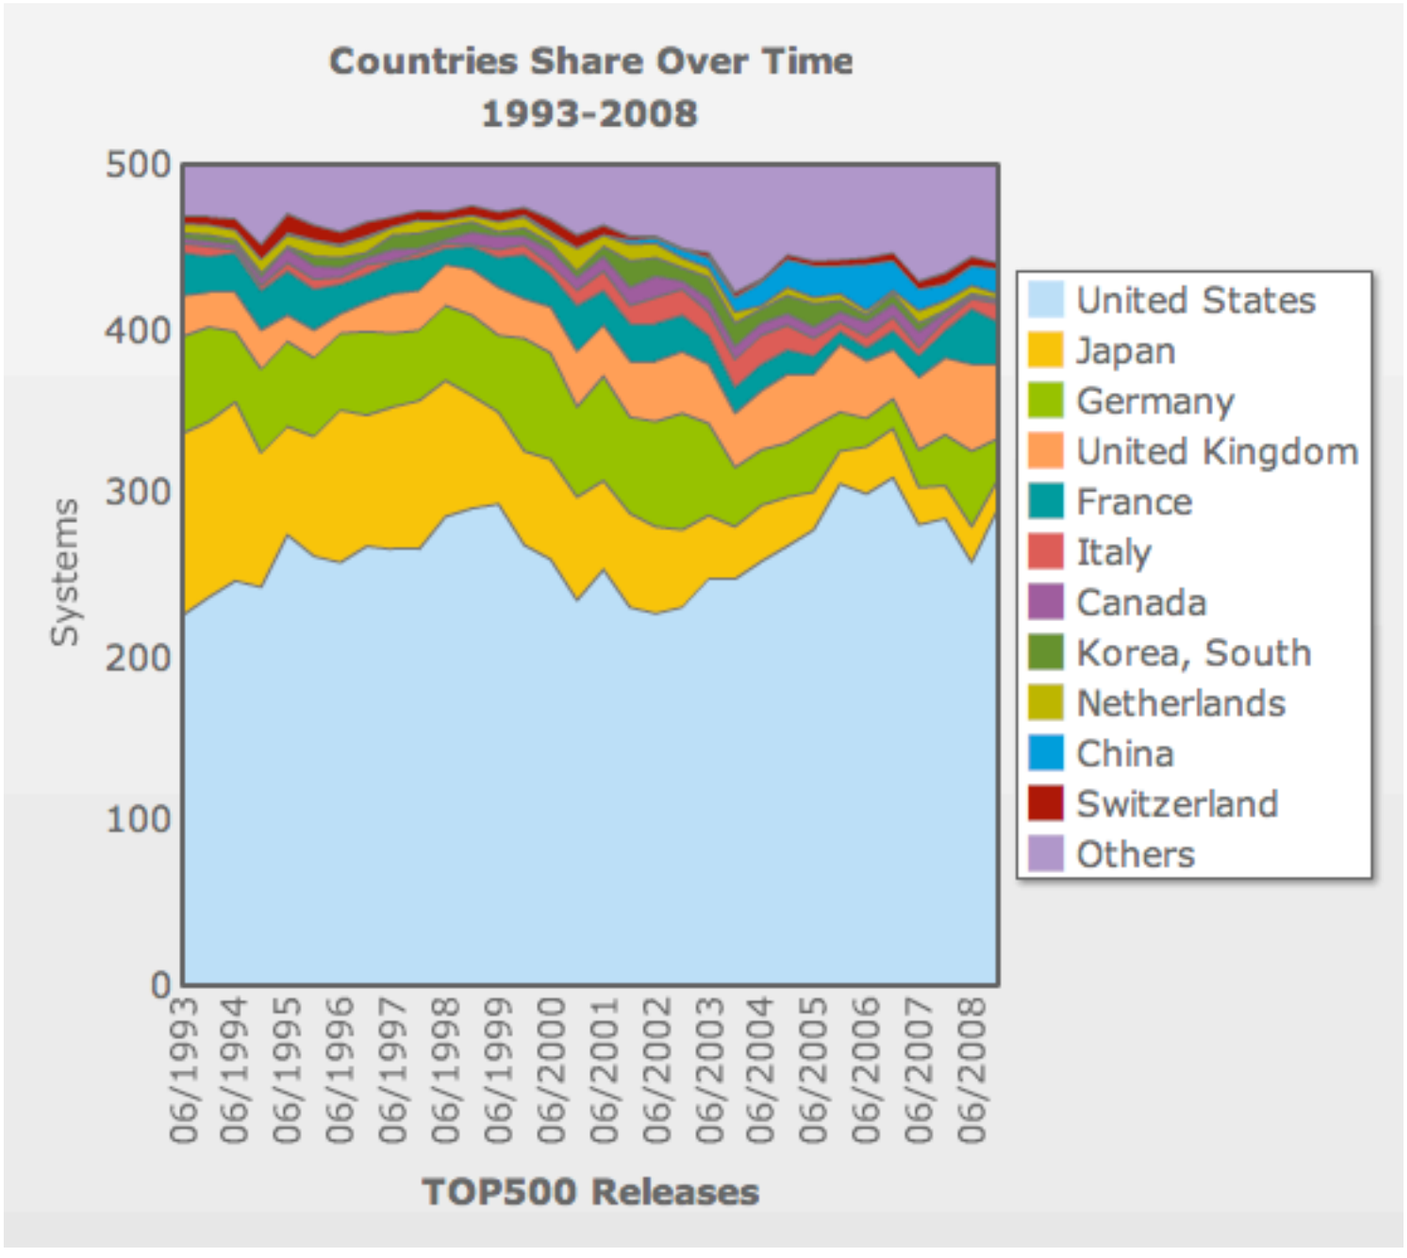

Evolution: country (top500)

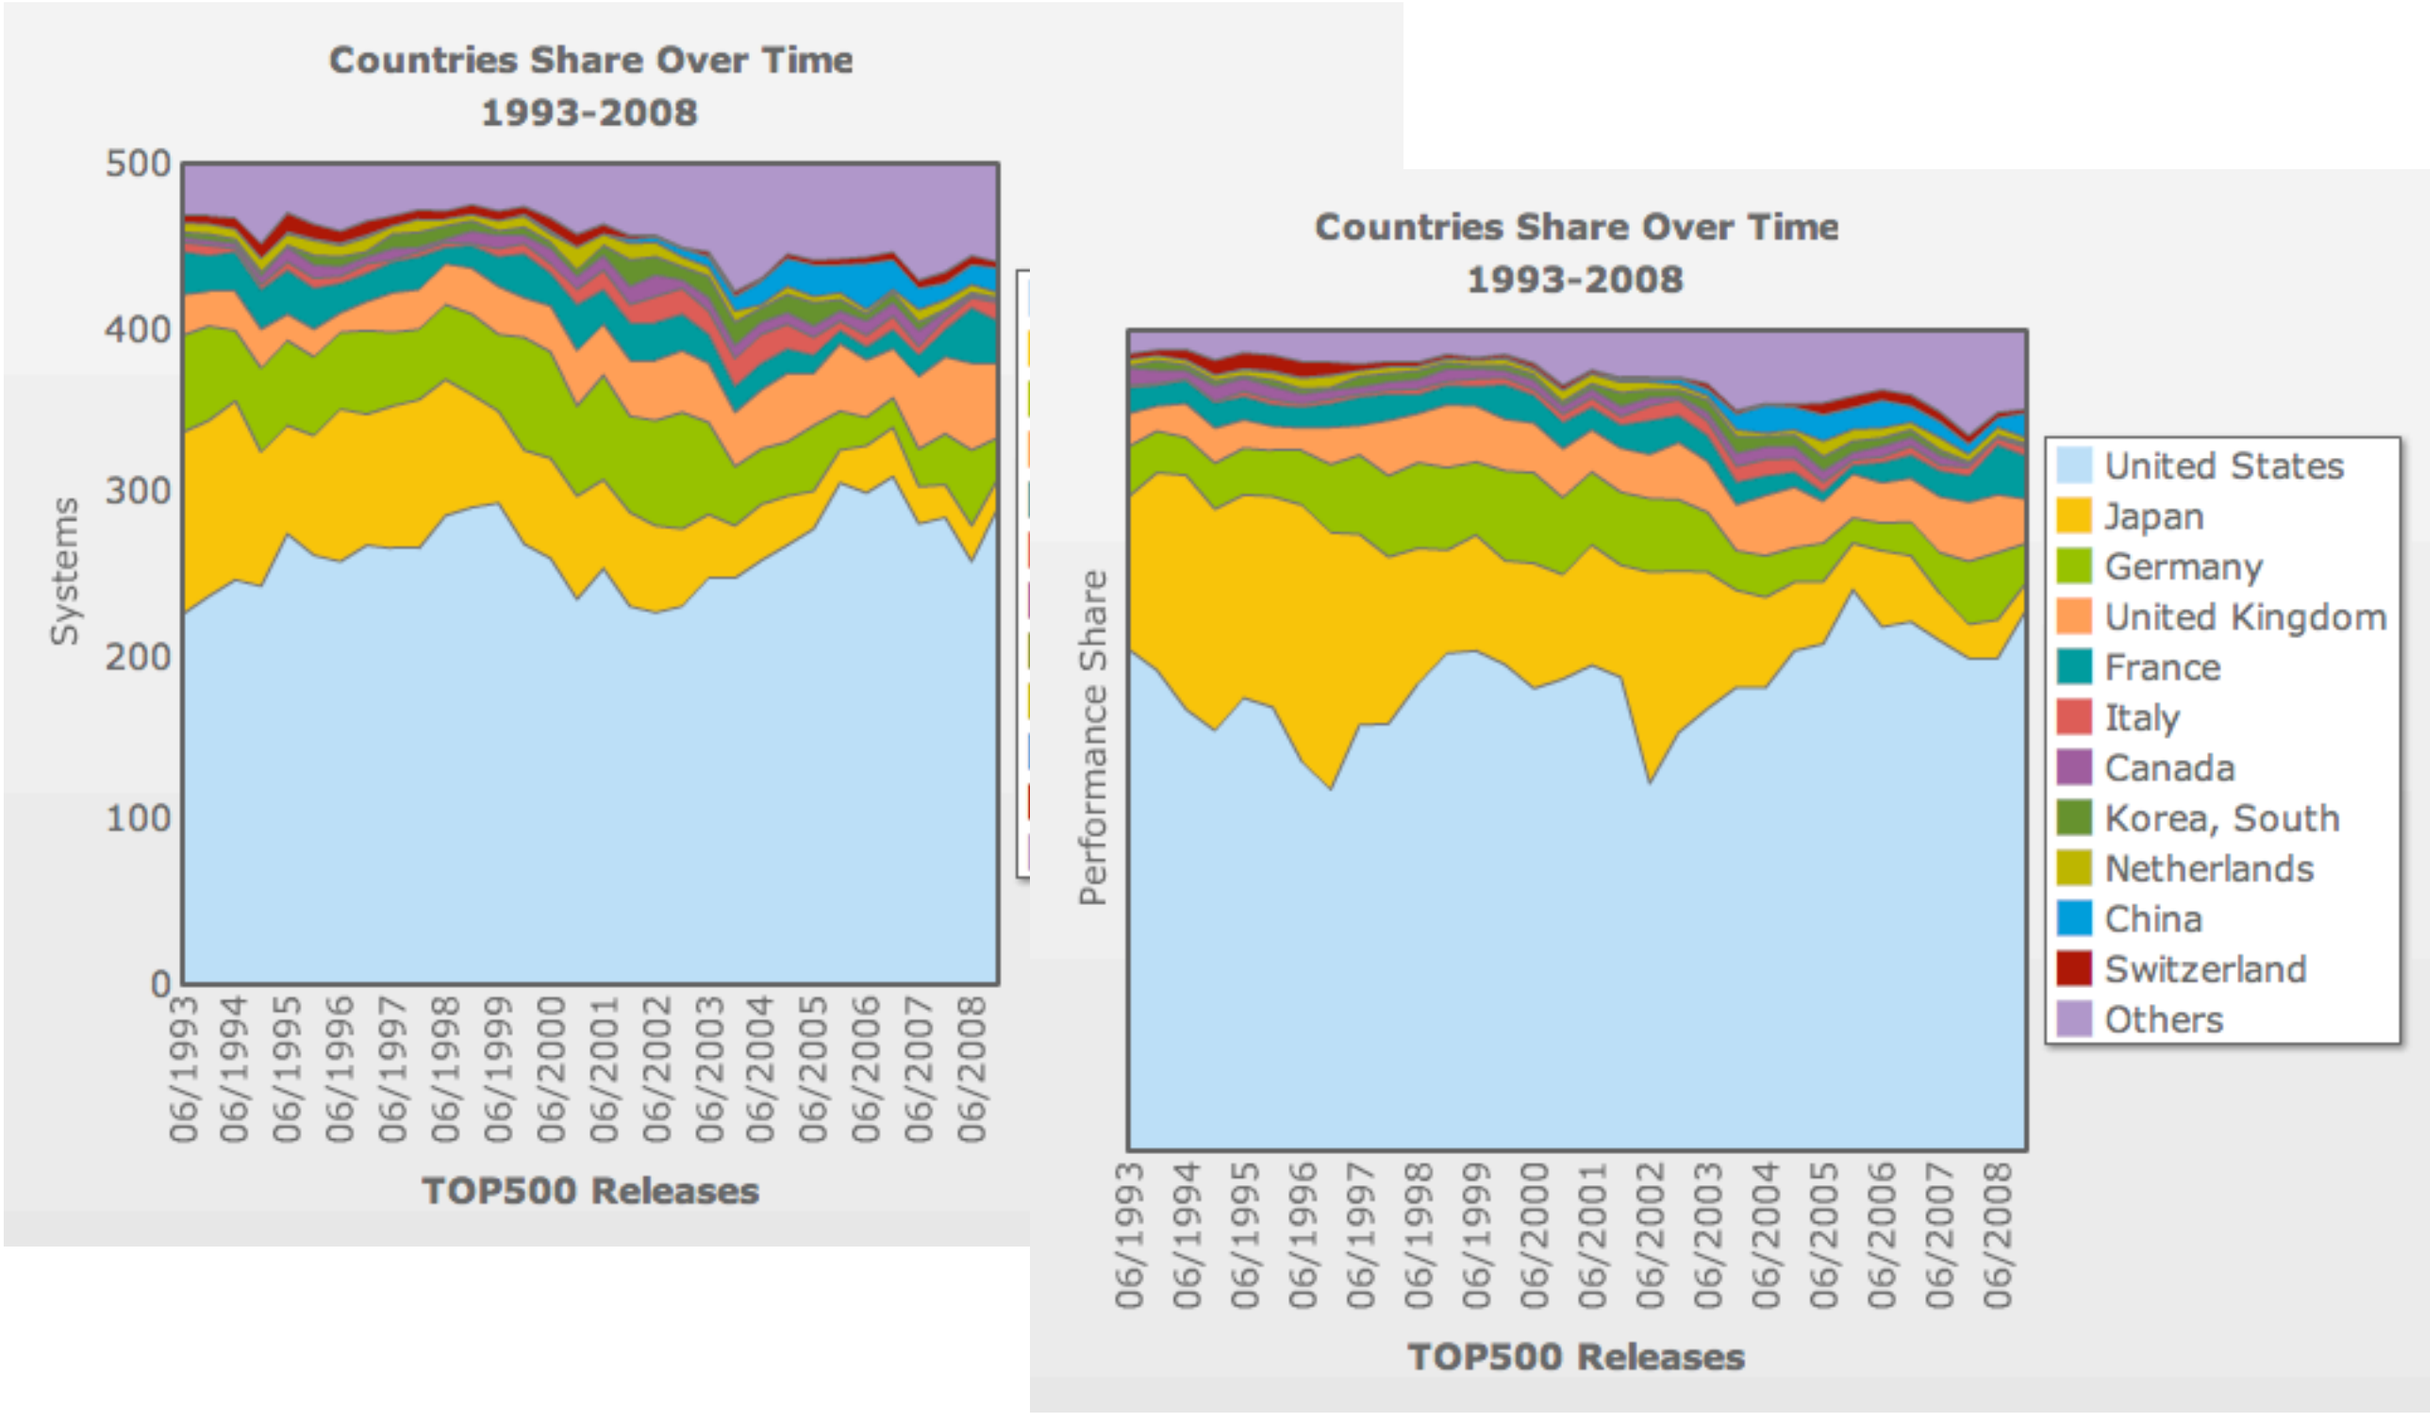

Evolution: applications

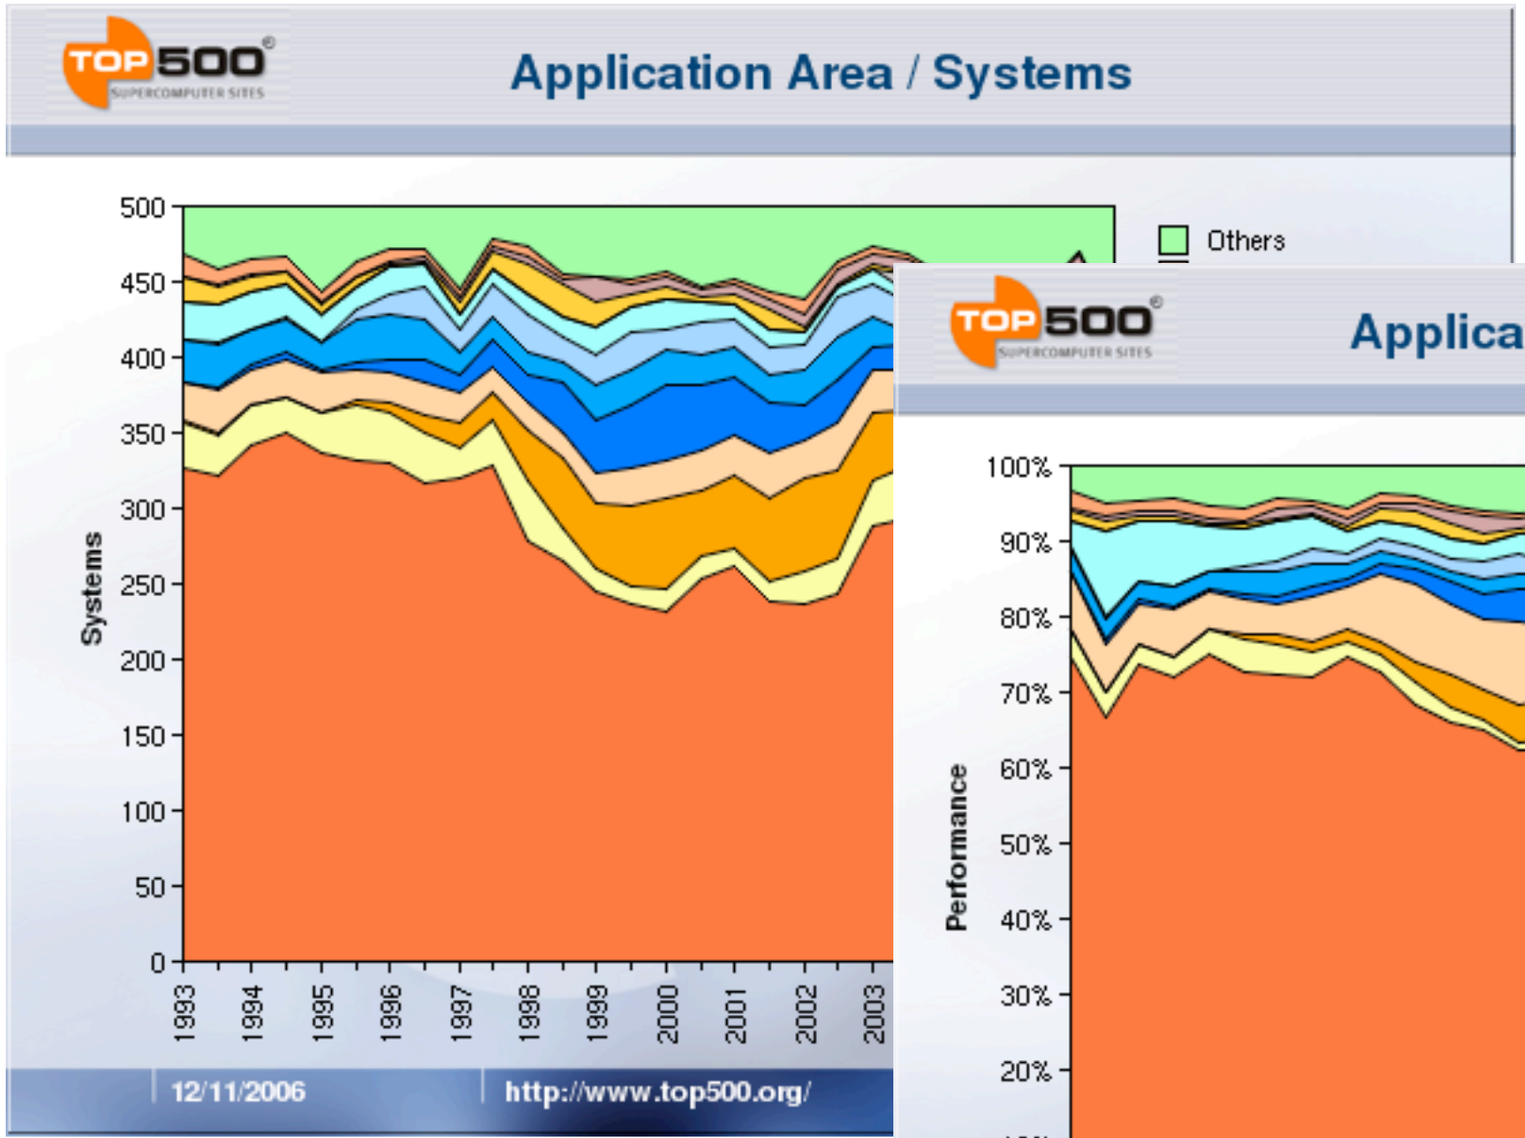

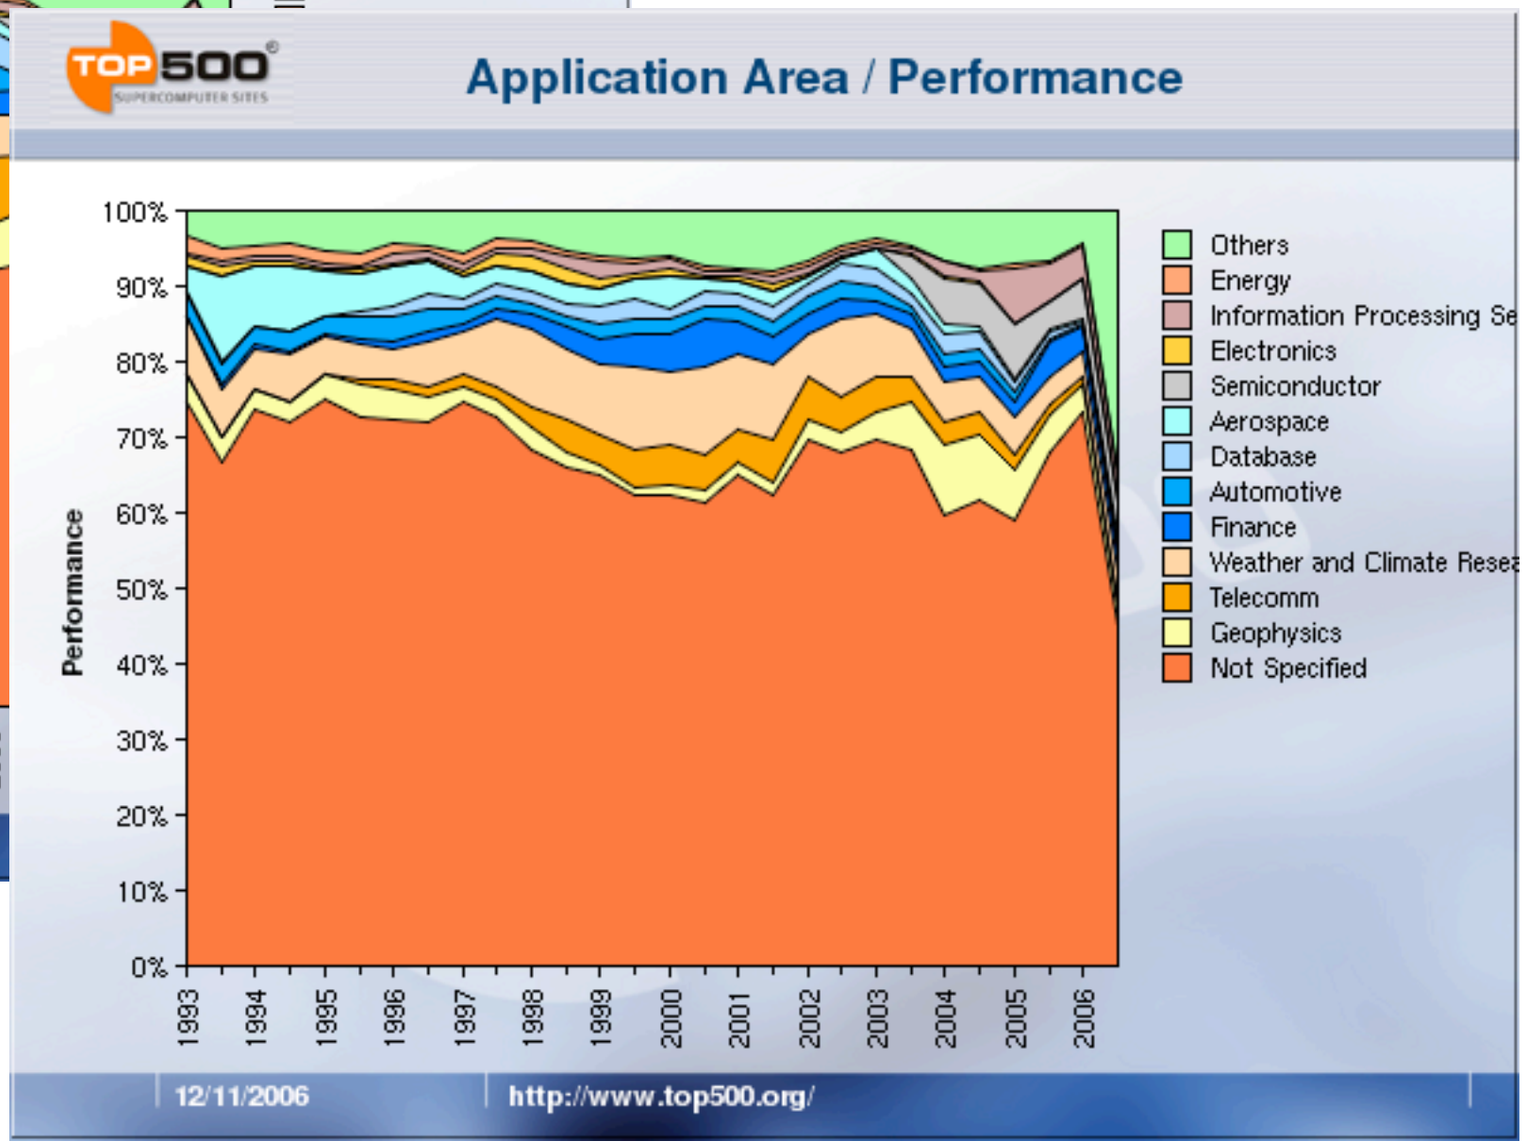

current Top500

TOP 10 Sites for November 2008

Development Over Time

For more information about the sites and systems in the list, click on the links or view the complete list.

Rank	Site	Computer
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband IBM
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT5 QC 2.3 GHz Cray Inc.
3	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz SGI
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution IBM
5	Argonne National Laboratory United States	Blue Gene/P Solution IBM
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband Sun Microsystems
7	NERSC/LBNL United States	Franklin - Cray XT4 QuadCore 2.3 GHz Cray Inc.
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz Cray Inc.
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core Cray Inc.
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008 Dawning

current Top500

TOP 10 Sites for November 2008

For more information about the sites and systems in or view the complete list.

Rank	Site	Computer
1	DOE/NNSA/LANL United States	Roadrunner DC 1.8 GHz IBM
2	Oak Ridge National Laboratory United States	Jaguar - Cra Cray Inc.
3	NASA/Ames Research Center/NAS United States	Pleiades - S SGI
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution IBM
5	Argonne National Laboratory United States	Blue Gene/P Solution IBM
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband Sun Microsystems
7	NERSC/LBNL United States	Franklin - Cray XT4 QuadCore 2.3 GHz Cray Inc.
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz Cray Inc.
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core Cray Inc.
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008 Dawning

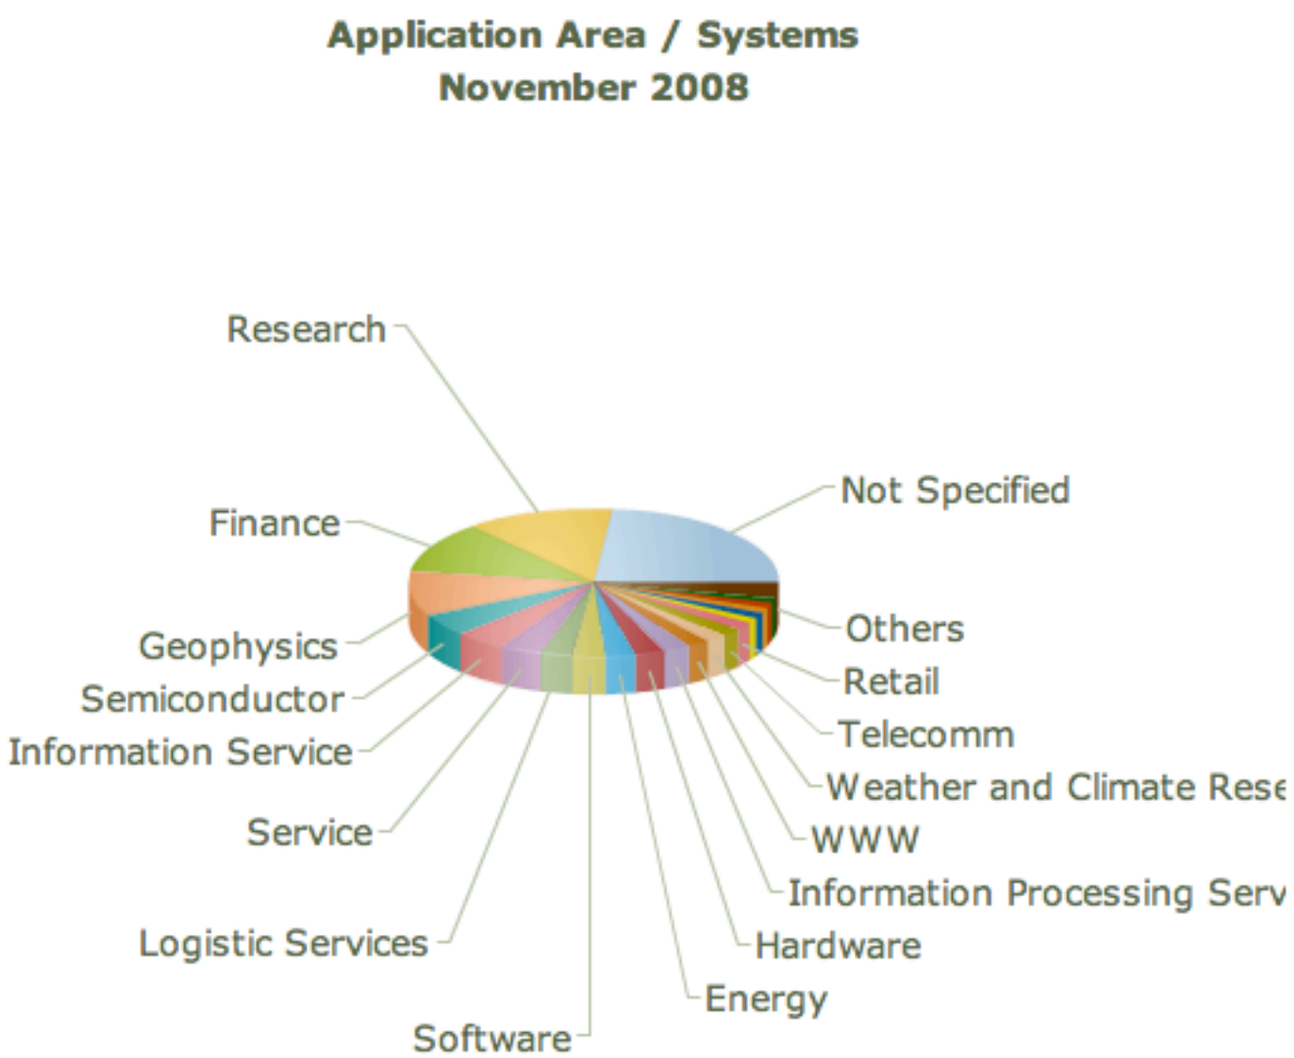

current Top500

TOP 10 Sites for November 2008

For more information about the sites and systems in
or view the complete list.

Rank	Site	Computer
1	DOE/NNSA/LANL United States	Roadrunner DC 1.8 GHz IBM
2	Oak Ridge National Laboratory United States	Jaguar - Cra Cray Inc.
3	NASA/Ames Research Center/NAS United States	Pleiades - S SGI
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blu IBM
5	Argonne National Laboratory United States	Blue Gene/P Solution IBM
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x642 Sun Microsystems
7	NERSC/LBNL United States	Franklin - Cray XT4 Quad Cray Inc.
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz Cray Inc.
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core Cray Inc.
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008 Dawning

Ir

Application Area / Systems
November 2008

Processor Architecture / Systems
November 2008

current Top500

TOP 10 Sites for November 2008

For more information about the sites and systems in
or view the complete list.

Rank	Site	Computer
1	DOE/NNSA/LANL United States	Roadrunner DC 1.8 GHz IBM
2	Oak Ridge National Laboratory United States	Jaguar - Cra Cray Inc.
3	NASA/Ames Research Center/NAS United States	Pleiades - S SGI
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blu IBM
5	Argonne National Laboratory United States	Blue Gene/P Solution IBM
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x642 Sun Microsystems
7	NERSC/LBNL United States	Franklin - Cray XT4 Quad Cray Inc.
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 Cray Inc.
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Sto Cray Inc.
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A 2008 Dawning

Ir

Application Area / Systems
November 2008

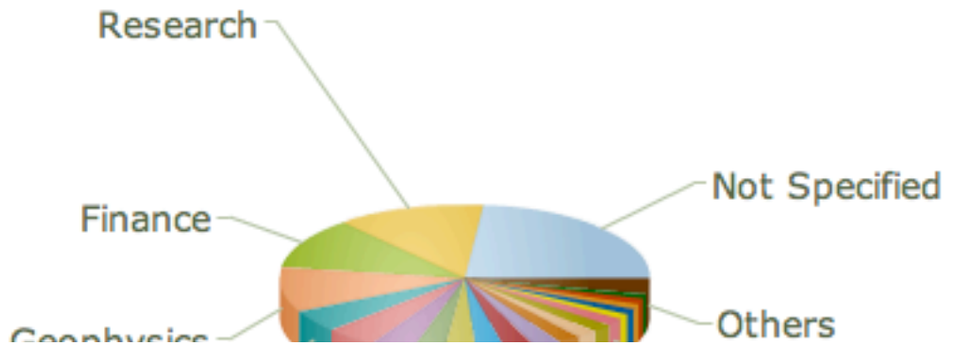

Processor Architecture / Systems
November 2008

Number of Processors / Systems
November 2008

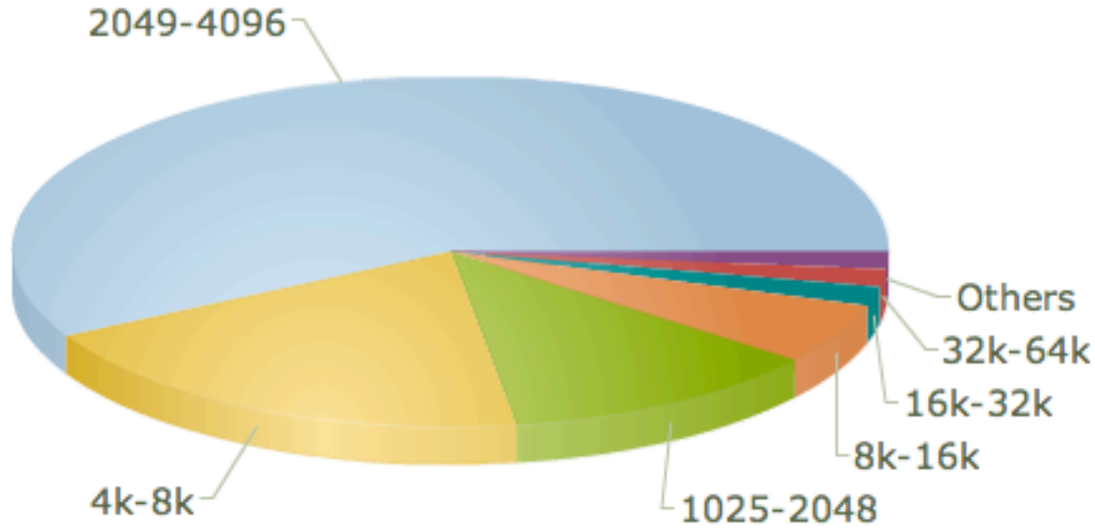

IBM Roadrunner

Roadrunner

Details | Performance/Linpack Data | Ranking History

System Name	**Roadrunner**
Site	DOE/NNSA/LANL
System Family	IBM Cluster
System Model	BladeCenter QS22 Cluster
Computer	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband
Vendor	IBM
Application area	Not Specified
Installation Year	2008
Operating System	Linux
Interconnect	Infiniband
Processor	PowerXCell 8i 3200 MHz (12.8 GFlops)

IBM Roadrunner

Roadrunner

Details

System Nam
Site
System Fami
System Mode
Computer
Vendor
Application ar
Installation Y

Operating Sy
Interconnect
Processor

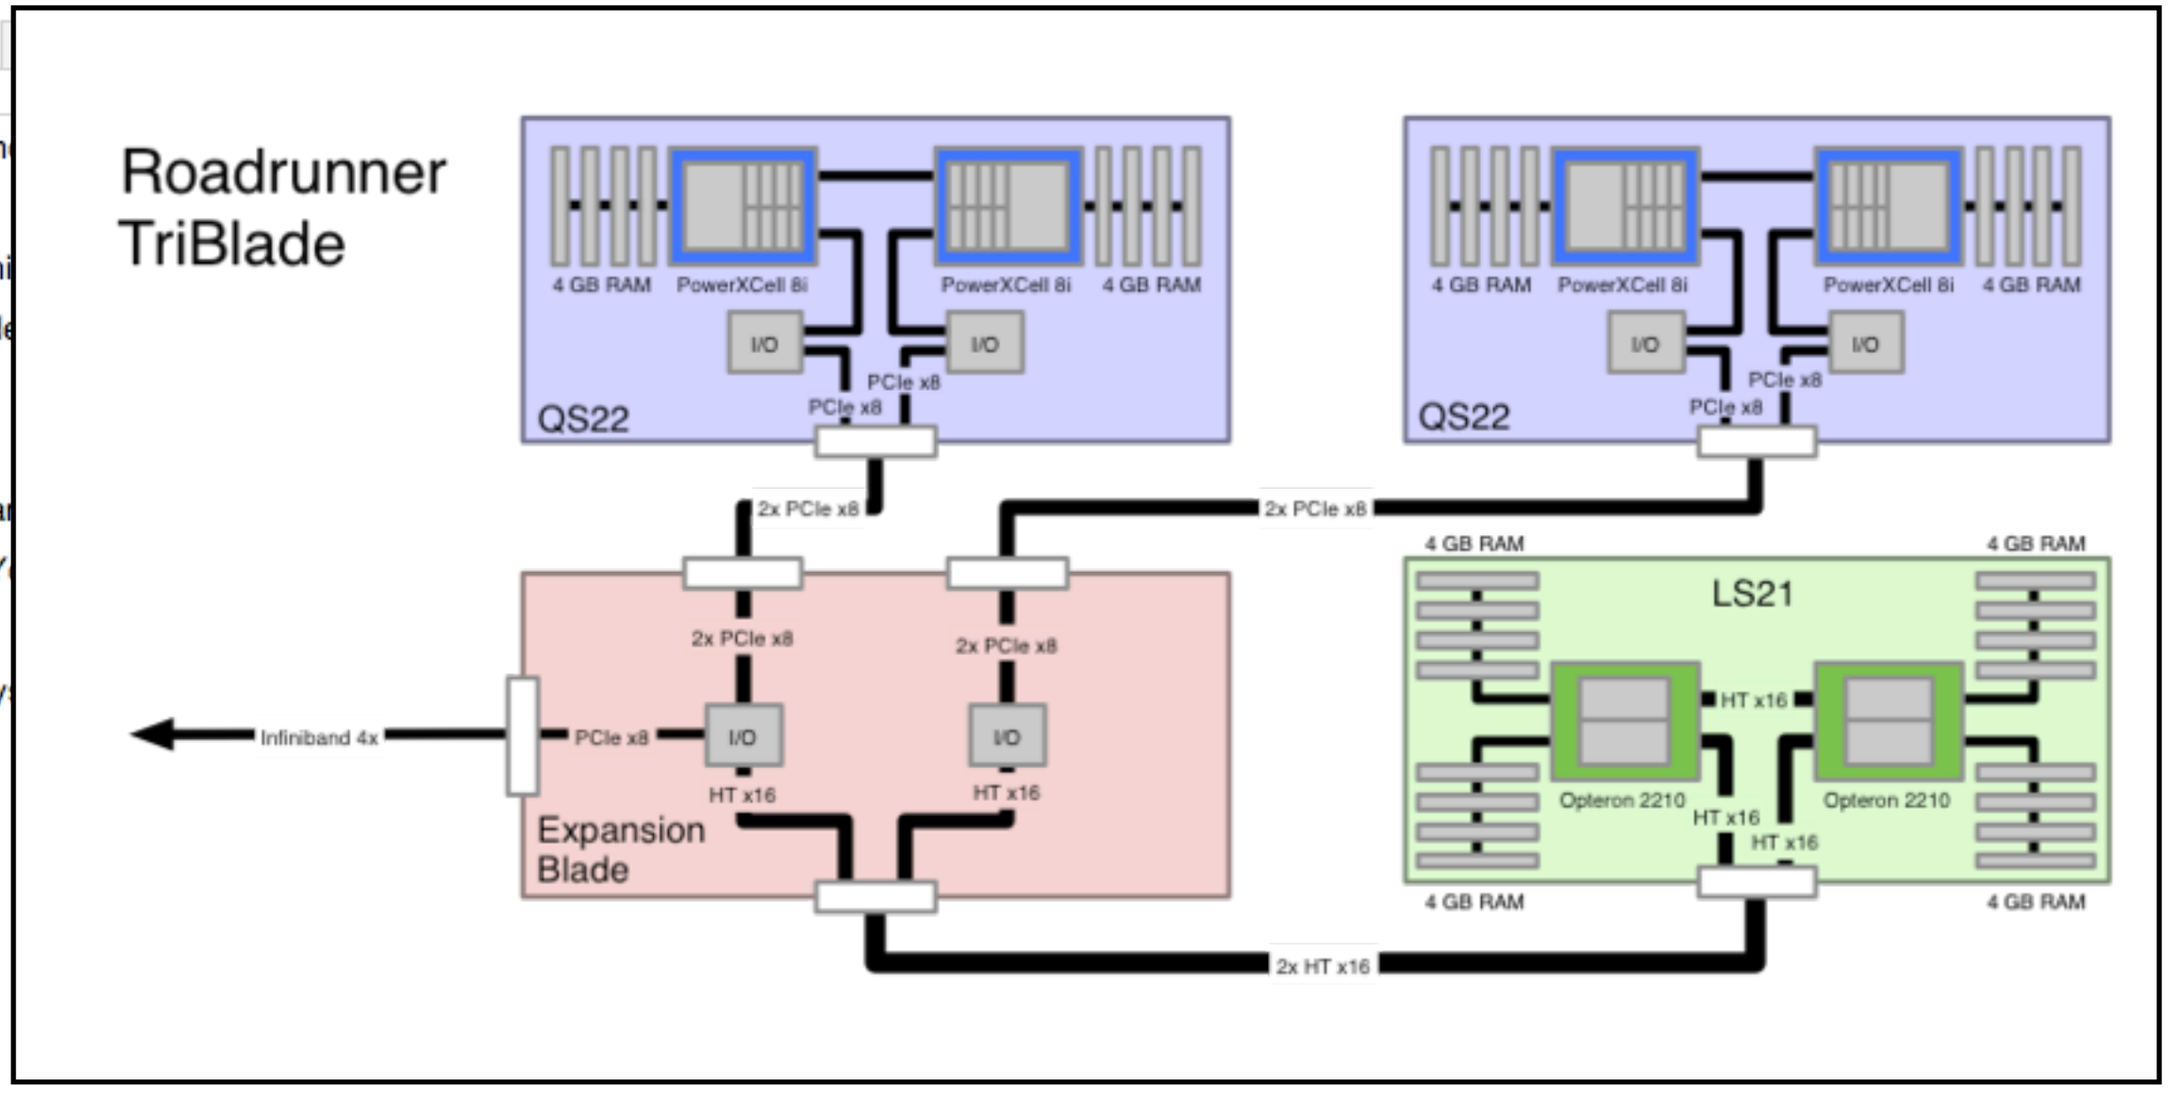

IBM Roadrunner

Roadrunner

Details

System Nam
Site
System Fami
System Mode
Computer
Vendor
Application ar

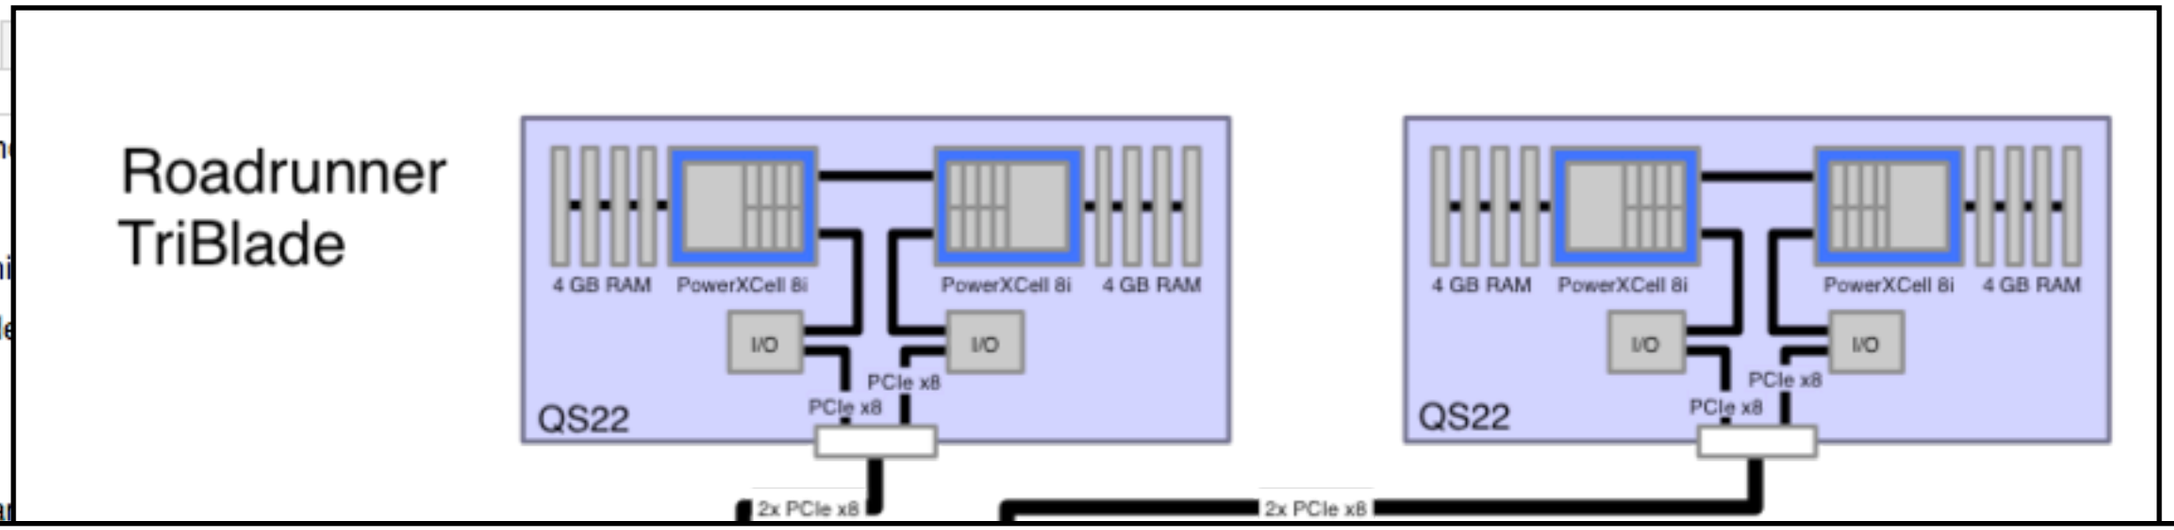

Roadrunner, tiered architecture

TriBlade

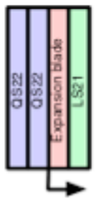

TriBlade:
4x PowerXCell 8i
2x Opteron 2210
32x GB RAM
~400 Gflops

BladeCenter H

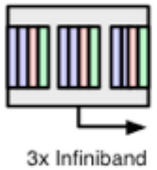

BladeCenter H:
3x TriBlade
12x PowerXCell 8i
6x Opteron 2210
96 GB RAM
~1.2 Tflops

Standard 42U Rack

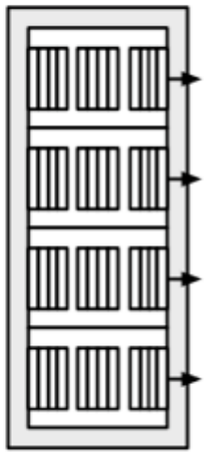

Rack:
4x BladeCenter H
12x TriBlades
48x PowerXCell 8i
24x Opteron 2210
384 GB RAM
~4.8 Tflops

Connected Unit

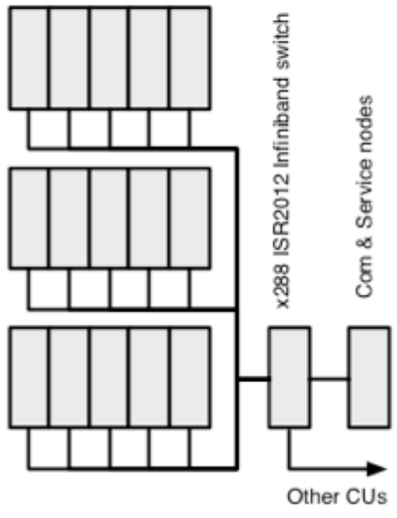

Connected Unit:
15x Racks
60x BladeCenter H
180x TriBlades
720x PowerXCell 8i
360x Opteron 2210
5.76 TB RAM
~72 Tflops

Com & Service node:
12x I/O nodes, IBM x3655
~250 Gflops
Service node, IBM x3655

Infiniband switch, Voltaire ISR2012

Roadrunner cluster

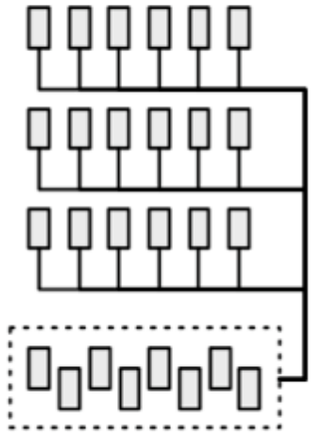

Roadrunner Cluster:
18x Connected Units
270x Racks
1080x BladeCenter H
3240x TriBlades
12960x PowerXCell 8i
6480x Opteron 2210
103.6 TB RAM
~1.3 Pflops

18x Com & Service nodes
~4.5 Tflops
18x Infiniband switches

2nd stage Infiniband:
x8 Infiniband switches

296x Racks in total

What's up in Italy ?

Countries share for 11/2008

In addition to the table below, you can view the visual charts using the TOP500 charts page. A direct link to the charts is also available.

Countries	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Australia	1	0.20 %	21906	40960	4096
Belgium	2	0.40 %	35432	58165	5552
Brazil	2	0.40 %	31778	53355	5008
Bulgaria	1	0.20 %	23415	27850	8192
Canada	2	0.40 %	70479	88954	6800
China	15	3.00 %	520962	844192	90400
Denmark	3	0.60 %	46472	55585	5440
Finland	1	0.20 %	76510	99507	10816
France	26	5.20 %	886782	1298890	184276
Germany	25	5.00 %	803761	1087557	169844
India	8	1.60 %	259394	368501	37488
Ireland	1	0.20 %	25110	28672	2560
Israel	1	0.20 %	24670	61152	5096
Italy	11	2.20 %	182434	307440	35248
Japan	17	3.40 %	522981	734926	106984
Korea, South	1	0.20 %	16990	24576	3072
Malaysia	1	0.20 %	20442	38224	4096
Mexico	1	0.20 %	18480	25440	2120
Netherlands	3	0.60 %	111506	138752	27904
New Zealand	4	0.80 %	109109	204800	20480
Norway	2	0.40 %	72454	110984	11182
Poland	6	1.20 %	116999	172675	16600
Russia	8	1.60 %	222968	285754	31840
Slovenia	1	0.20 %	36752	49152	4096
South Africa	1	0.20 %	23415	27850	8192
Spain	6	1.20 %	153201	231953	26616
Sweden	8	1.60 %	274791	405047	39084

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

- 2000
- 2006

“Desk processor”

- http://www.intel.com/pres
- 1993
- 2000
- 2006

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

Desktop				
Processor	Clock Speed(s)	Intro Date(s)	Mfg. Process/ Transistors	Typical Use
Intel® Pentium® Processor with MMX™ Technology	233 MHz	June 2, 1997	0.35-micron 4.5 million	High-performance desktops and servers
Intel® Pentium® Processor with MMX™ Technology	200 MHz 166 MHz	Oct. xx, 1996	0.35-micron 4.5 million	High-performance desktops and servers
Intel® Pentium® Processor	200 MHz	June 10, 1996	0.35-micron 3.3 million	High-performance desktops and servers
Intel® Pentium® Processor	166 MHz 150 MHz	Jan. 4, 1996	0.35-micron 3.3 million	High-performance desktops and servers
Intel® Pentium® Processor	133 MHz	June 1995	0.35-micron 3.3 million	High-performance desktops and servers
Intel® Pentium® Processor	120 MHz	Mar. 27, 1995	0.6-micron 0.35-micron 3.3 million	Desktops and notebooks
Intel® Pentium® Processor	100 MHz 90 MHz	Mar. 7, 1994	0.6-micron 3.3 million	Desktops
Intel® Pentium® Processor	75 MHz	Oct. 10, 1994	0.6-micron 3.3 million	Desktops and notebooks
Intel® Pentium® Processor	66 MHz 60 MHz	Mar. 22, 1993	0.8-micron 3.1 million	Desktops

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

- 2000
- 2006

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

February 14, 2000
Mobile Intel® Celeron® Processor
500 MHz, 450 MHz

- 2000

January 18, 2000
Mobile Intel® Pentium® III Processor
650 MHz, 600 MHz

January 12, 2000
Intel® Pentium® III Xeon™ Processor
800 MHz

- 2006

“Desk processor”

- http://www.intel.com/pres
- 1993
- 2000
- 2006

1993

March 22, 1993
Intel® Pentium® Processo
66 MHz, 60 MHz

February 14, 2000
Mobile Intel® Celeron® Pr
500 MHz, 450 MHz

January 18, 2000
Mobile Intel® Pentium® II
650 MHz, 600 MHz

January 12, 2000
Intel® Pentium® III Xeon™
800 MHz

Intel® Pentium® III Xeon™ Processor

Processor	Clock Speed(s)	Intro Date(s)	Mfg. Process/ Transistors	Cache	Addressable Memory	Bus Speed	Typical Use
Intel® Pentium® III Xeon™ Processor	900 MHz	Mar. 21, 2001	0.18-micron 28 million	2 MB Advanced Transfer L2 cache	64 GB	100 MHz	High-end servers, 4- and 8-way multiprocessing systems
Intel® Pentium® III Xeon™ Processor	933 MHz	May 24, 2000	0.18-micron 28 million	256 KB Advanced Transfer L2 cache	64 GB	133 MHz	Business and consumer PCs, 1- and 2-way servers and workstations
Intel® Pentium® III Xeon™ Processor	700 MHz	May 22, 2000	0.18-micron 28 million	1 MB and 2 MB Advanced Transfer L2 cache	64 GB	100 MHz	4- and 8-way servers
Intel® Pentium® III Xeon™ Processor	866 MHz 800 MHz 733 MHz 667 MHz 600 MHz	Apr. 10, 2000 866 MHz Jan. 12, 2000 800 MHz Oct. 25, 1999 733 MHz 667 MHz 600 MHz	0.18-micron 28 million	256 KB Advanced Transfer L2 cache	64 GB	133 MHz	2-way servers and workstations
Intel® Pentium® III Xeon™ Processor	550 MHz 500 MHz	Mar. 17, 1999	0.25-micron 9.5 million	512 KB, 1 MB and 2 MB Advanced Transfer L2 cache	64 GB	100 MHz	Business PCs, 2-, 4- and 8-way (and higher) servers and workstations

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

February 14, 2000
Mobile Intel® Celeron® Processor
500 MHz, 450 MHz

- 2000

January 18, 2000
Mobile Intel® Pentium® III Processor
650 MHz, 600 MHz

January 12, 2000
Intel® Pentium® III Xeon™ Processor
800 MHz

- 2006

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

February 14, 2000
Mobile Intel® Celeron® Processor
500 MHz, 450 MHz

- 2000

January 18, 2000
Mobile Intel® Pentium® III Processor
650 MHz, 600 MHz

January 12, 2000
Intel® Pentium® III Xeon™ Processor
800 MHz

- 2006

Intel® Core™2 Duo Processor

Great computing starts with Intel inside

“Desk processor”

- http://www.intel.com/pressroom/kits/quickrefyr.htm
- 1993
- 2000
- 2006

1993

March 22, 1993
Intel® Pentium® Processor
66 MHz, 60 MHz

February 14,
Mobile Intel®
500 MHz, 450

January 18, 2
Mobile Intel®
650 MHz, 600

January 12, 2
Intel® Pentiun
800 MHz

Intel® Co
Great computin

	Intel® Itanium® 2 Processor	Intel® Itanium® 2 Processor Low Voltage
	Remove	Remove
Clock Speed	1.66 GHz	1.30 GHz
Front Side Bus Speed	667 MHz	400 MHz
L3 Cache	9MB	3MB
L2 Cache	256KB	256KB
L1 Cache	32KB	32KB
Power	130 watts	62 watts
System Type	MP	DP
Architecture	.13 micron	.13 micron
Other Intel Technologies	EPIC, Machine Check Architecture	EPIC, Machine Check Architecture
Package	PAC-611	PAC-611
Server Chipset	Intel® E8870 Chipset, OEM custom chipset	Intel® E8870 Chipset, OEM custom chipset
Memory Type	DDR, SDRAM	DDR, SDRAM
Server Platforms	Intel® Server Platform SR870BN4	Intel® Server Platform SR870BH2

Currently

Intel® Core™ Duo processor

Intel's first mobile dual-core processor executes multiple threads simultaneously using two cores, thereby helping to maximize performance and multitasking capabilities.

» **Learn more**

Processor Number[1]	Architecture	Cache	Clock Speed	Front Side Bus	Power	Dual-core	Intel® Virtualization Technology (Intel® VT)Φ	Enhanced Intel SpeedStep® Technology	Execute Disable Bit°
T2700	65 nm	2 MB L2	2.33 GHz	667 MHz	31W	✓	✓	✓	
T2600	65 nm	2 MB L2	2.16 GHz	667 MHz	31W	✓	✓	✓	✓
T2500	65 nm	2 MB L2	2 GHz	667 MHz	31W	✓	✓	✓	✓
T2450	65 nm	2 MB L2	2 GHz	533 MHz	31W	✓		✓	✓
T2400	65 nm	2 MB L2	1.83 GHz	667 MHz	31W	✓	✓	✓	✓
T2350	65 nm	2 MB L2	1.86 GHz	533 MHz	31W	✓		✓	✓
T2300	65 nm	2 MB L2	1.66 GHz	667 MHz	31W	✓	✓	✓	✓
T2250	65 nm	2 MB L2	1.73 GHz	533 MHz	31W	✓		✓	✓
T2050	65 nm	2 MB L2	1.6 GHz	533 MHz	31W	✓		✓	✓
T2300E	65 nm	2 MB L2	1.66 GHz	667 MHz	31W	✓		✓	✓

Currently

Intel® Core™ Duo processor

Intel's first mobile dual-core processor executes multiple threads simultaneously using two cores, thereby helping to maximize perfc

» Learn more

Processor Number[1]	Architecture	
T2700	65 nm	2
T2600	65 nm	2
T2500	65 nm	2
T2450	65 nm	2
T2400	65 nm	2
T2350	65 nm	2
T2300	65 nm	2
T2250	65 nm	2
T2050	65 nm	2
T2300E	65 nm	2

Featuring
Intel® Core™ Duo processor low voltage

Processor Number[1]	Architecture	Cache	Clock Speed	Front Side Bus	Power	Dual-core	Intel® Virtualization Technology (Intel® VT)Φ	Enhanced Intel SpeedStep® Technology	Execute Disable Bit°
L2500	65 nm	2MB L2	1.83 GHz	667 MHz	15W	✓	✓	✓	✓
L2400	65 nm	2MB L2	1.66 GHz	667 MHz	15W	✓	✓	✓	✓
L2300	65 nm	2MB L2	1.50 GHz	667 MHz	15W	✓	✓	✓	✓

Featuring
Intel® Core™ Duo processor ultra low voltage

Processor Number[1]	Architecture	Cache	Clock Speed	Front Side Bus	Power	Dual-core	Intel® Virtualization Technology (Intel® VT)Φ	Enhanced Intel SpeedStep® Technology	Execute Disable Bit°
U2500	65 nm	2MB L2	1.20 GHz	533 MHz	9W	✓	✓	✓	✓
U2400	65 nm	2MB L2	1.06 GHz	533 MHz	9W	✓	✓	✓	✓

Currently

Intel® Core™ Duo processor

Intel's first mobile dual-core processor executes multiple threads simultaneously using two cores, thereby helping to maximize perfc

Featuring

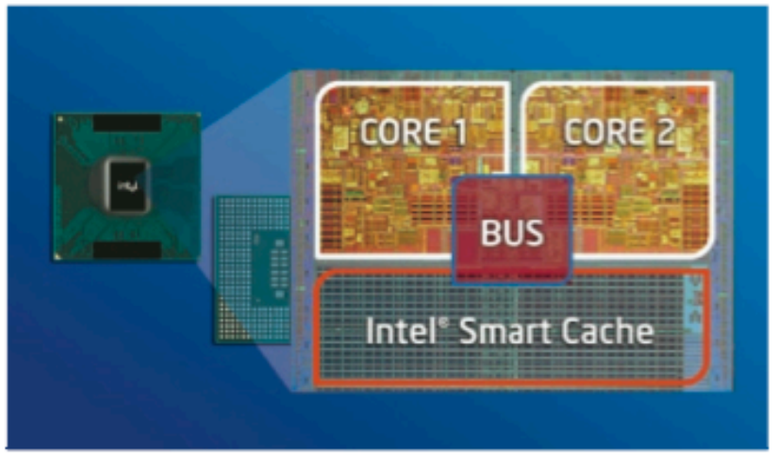

Specialized Dual-Core Microarchitecture

The Intel Core Duo processor includes two mobile-optimized execution cores in a single processor. This design enables execution of parallel threads or applications on separate cores with dedicated CPU resources. The results enable outstanding dual-core performance and greater system response when running multi-threaded or multiple demanding applications simultaneously.†

The Intel Core Duo processor features a high-performance

Features and Benefits of the Intel® Core™ Duo Processor

Features	Benefits
Intel® Smart Cache – 2MB L2 cache with Advanced Transfer Cache Architecture	Delivers a smarter and more efficient cache and bus design to enable enhanced dual-core performance and power savings.
Intel® Digital Media Boost	Micro-architectural enhancements that include instruction optimizations and performance enhancements accelerate a diverse variety of processing-intensive tasks, such as audio/video processing, image processing, 3D graphics, and scientific calculations.
Intel® Dynamic Power Coordination with Dynamic Bus Parking	Dual-core on demand, coordinated performance with enhanced low power management features Dynamic Bus Parking. This enables platform power savings by allowing the chipset to power down with the processor in these low-frequency mode states.
Enhanced Intel® Deeper Sleep with Dynamic Cache Sizing	Allows the processor to lower voltage below the Deeper Sleep minimum voltage to enable enhanced power savings. Dynamic Cache Sizing is a new power savings mechanism that enables the Intel® Smart Cache to dynamically flush system memory based on demand or during periods of inactivity.
Intel® Advanced Thermal Manager	A new thermal management system delivers enhanced accuracy and more precise acoustic control to enable quieter, cooler, thinner system designs.
Power-Optimized 667 MHz System Bus	Utilizes Source-Synchronous Transfer (SST) of address and data enables improved performance by transferring data at 4X bus clock. Advanced Gunning Transceiver Logic (AGTL+) signaling technology, a variant of GTL+ signaling technology, delivering low power enhancements.
Enhanced Intel SpeedStep® Technology Support	Multiple performance modes enable optimum performance at the lowest power, using real-time dynamic switching of the voltage and frequency between multiple performance modes based on CPU demand.
New Intel 65nm Process Technology	Smaller transistors enable more logic and more frequency headroom for increased performance.

Processor Number[1]	Architecture	Cache	Clock Speed	Side Bus	Power	Dual-core	Virtualization Technology (Intel® VT)Φ	Intel SpeedStep® Technology	Execute Disable Bit°
U2500	65 nm	2MB L2	1.20 GHz	533 MHz	9W	✓	✓	✓	✓
U2400	65 nm	2MB L2	1.06 GHz	533 MHz	9W	✓	✓	✓	✓

Currently (2)

Access. Get unmatched quad-core desktop performance[1] for intensive multimedia applications such as video compression, photo editing, retouching, and publishing. With Intel® HD Boost, you will experience higher performance for intensive applications such as video compression while maintaining high visual quality.

Multitasking Monster

The Intel® Core™2 Extreme Quad-Core processor is a multitasking monster, so users can do more in less time. Significant improvements in system responsiveness are possible because certain tasks can be off-loaded to specific cores. Users can now take on several tasks at once, such as rendering a video, playing a game, or working on basic productivity software, because additional processor resources are free to handle other tasks.

Better Acoustics

Intel Core 2 Extreme Quad-Core processors are equipped with a Digital Thermal Sensor (DTS) to enable more efficient processor and platform thermal control. Thermal sensors located within the processor measure the maximum temperature on the die at any given time. The acoustic benefit of temperature monitoring is that system fans spin only as fast as needed to cool the system and slower spinning fans generate less noise.

This Intel-designed thermal solution for boxed processors utilizes a 4-pin header with variable fan speed control, based on processor temperature and power usage to minimize acoustic noise levels. The latest 45nm boxed Intel Core 2 Extreme processors include a new, more advanced thermal solution with improved acoustic and thermal performance.

45nm Comparison Table

	QX9770	QX9650
Clock Speed	3.2 GHz	3 GHz
L2 Cache[2]	12 MB	12 MB
Front Side Bus Speed	1600 MHz	1333 MHz
Intel® Express Chipset	X48	X48, X38, P35

Features and Benefits of the Intel® Core™2 Extreme Quad-Core Processor

Feature	Benefit
Quad-Core Processing	Provides four independent execution cores in a single processor package. Four dedicated processing threads help operating systems and applications deliver additional performance, so end users can experience better multitasking and multithreaded performance across many types of applications and workloads.
Chipset Support	Intel® Express Chipsets offer an array of exciting capabilities including dual graphics, and deliver an impressive level of performance for demanding users. Other third-party chipsets may support Intel® Core™2 Extreme processors; contact your board manufacturer for compatibility.
Intel® Wide Dynamic Execution	Improves execution speed and efficiency, delivering more instructions per clock cycle. Each core can complete up to four full instructions simultaneously.
Intel® Smart Memory Access	Improves system performance by optimizing the use of the available data bandwidth from the memory subsystem and reducing the effective latency of memory accesses.
Intel® Advanced Smart Cache[2]	Dynamically allocates the shared L2 cache to each processor core based on workload. This efficient, dual-core-optimized implementation increases the probability that each core can access data from fast L2 cache, significantly reducing latency to frequently used data and improving performance.
Intel® HD Boost	Accelerates the execution of Streaming SIMD Extension (SSE) instructions to significantly improve the performance on a broad range of multimedia and compute-intensive applications. The 128-bit SSE instructions are now issued at a throughput rate of one per clock cycle, effectively doubling their speed of execution on a per-clock basis over previous generation processors. This is now improved further on 45nm versions with new SSE4 instructions for even better multimedia performance.
Intel® Virtualization Technology (Intel® VT)[3]	Allows one hardware platform to function as multiple "virtual" platforms. Intel VT improves manageability, limits downtime, and maintains worker productivity by isolating computing activities into separate partitions.
Intel® 64 Architecture[4]	Allows the processor to access larger amounts of memory. With appropriate 64-bit hardware and software, platforms based on an Intel® processor supporting Intel 64 architecture can allow the use of extended virtual and physical memory.
Execute Disable Bit[5]	Provides extended virus defense when deployed with a supported operating system. Memory can be marked as executable or non-executable, allowing the processor to raise an error to the operating system if malicious code attempts to run in non-executable memory. This can prevent the code from infecting the system.

Currently (2)

Access. Get unmatched quad-core desktop performance[1] for intensive multimedia applications such as video compression, photo editing, retouching, and publishing. With Intel® HD Boost, you will experience higher performance for intensive applications such as video compression while maintaining high visual quality.

Multitasking Monster

The Intel® Core™2 Extreme Quad-Core processor is a multitasking monster, so users can do more in less time. Significant improvements in system responsiveness are possible because certain tasks can be off-loaded to specific cores. Users can now take on several tasks at once, such as rendering a video, playing a game, or working on basic productivity software, because additional processor resources are free to handle other tasks.

Better Acoustics

Intel Core 2 Extreme Quad-Core processors are equipped with a Digital Thermal Sensor (DTS) to enable more efficient processor and platform thermal control. Thermal sensors located within the processor measure the maximum temperature on the die at any given time. The acoustic benefit of temperature monitoring is that system fans spin only as fast as needed to cool the system and slower spinning fans generate less noise.

This Intel-designed thermal solution for boxed processors utilizes a 4-pin header with variable fan speed control, based on processor temperature and power usage to minimize acoustic noise levels. The latest 45nm boxed Intel Core 2 Extreme processors include a new, more advanced thermal solution with improved acoustic and thermal performance.

45nm Comparison Table

	QX9770	QX9650
Clock Speed	3.2 GHz	3 GHz
L2 Cache[2]	12 MB	12 MB
Front Side Bus Speed	1600 MHz	1333 MHz
Intel® Express Chipset	X48	X48, X38, P35

Features and Benefits of the Intel® Core™2 Extreme Quad-Core Processor

Feature	Benefit
Quad-Core Processing	Provides four independent execution cores in a single processor package. Four dedicated processing threads help operating systems and applications deliver additional performance, so end users can experience better multitasking and multithreaded performance across many types of applications and workloads.
Chipset Support	Intel® Express Chipsets offer an array of exciting capabilities including dual graphics, and deliver an impressive level of performance for demanding users. Other third-party chipsets may support Intel® Core™2 Extreme processors; contact your board manufacturer for compatibility.
Intel® Wide Dynamic Execution	Improves execution speed and efficiency, delivering more instructions per clock cycle. Each core can complete up to four full instructions simultaneously.
Intel® Smart Memory Access	Improves system performance by optimizing the use of the available data bandwidth from the memory subsystem and reducing the effective latency of memory accesses.
Intel® Advanced Smart Cache[2]	Dynamically allocates the shared L2 cache to each processor core based on workload. This efficient, dual-core-optimized implementation increases the probability that each core can access data from fast L2 cache, significantly reducing latency to frequently used data and improving performance.
Intel® HD Boost	Accelerates the execution of Streaming SIMD Extension (SSE) instructions to significantly improve the performance on a broad range of multimedia and compute-intensive applications. The 128-bit SSE instructions are now issued at a throughput rate of one per clock cycle, effectively doubling their speed of execution on a per-clock basis over previous generation processors. This is now improved further on 45nm versions with new SSE4 instructions for even better multimedia performance.
Intel® Virtualization Technology (Intel® VT)[3]	Allows one hardware platform to function as multiple "virtual" platforms. Intel VT improves manageability, limits downtime, and maintains worker productivity by isolating computing activities into separate partitions.
Intel® 64 Architecture[4]	Allows the processor to access larger amounts of memory. With appropriate 64-bit hardware and software, platforms based on an Intel® processor supporting Intel 64 architecture can allow the use of extended virtual and physical memory.
Execute Disable Bit[5]	Provides extended virus defense when deployed with a supported operating system. Memory can be marked as executable or non-executable, allowing the processor to raise an error to the operating system if malicious code attempts to run in non-executable memory. This can prevent the code from infecting the system.

Solution for Boxed Processors	from running the fan at higher speeds for thermal performance.[3] Fan speed control technology is based on actual CPU temperature and power usage.

Then ?

- Proof of concept per 80-core !!!

Intel fabs 80-core teraflop processor

By Tony Smith in San Francisco [More by this author]
26th September 2006 20:44 GMT

Receive the days biggest stories by email, sign up here

IDF Quad-core? Pah! Intel has produced an 80-core chip, the world's first programmable microprocessor with teraflop performance capabilities, the chip giant claimed today. It's not compatible with the x86 instruction set - it's a proof of concept part designed to show how a production processor might operate.

The monster part incorporates not only the usual data-processing facilities - essentially they're just floating point maths co-processors - but also features a network processing unit on each core to control core-to-core communication. The cores are linked in a mesh configuration.

Each core's designed to be clocked to 3.1GHz and is mounted with 20MB of SRAM stacked up on top of the die. Connecting memory this way provides an aggregate bandwidth of a trillion bytes per second, Intel said.

80 core Intel proof of concept

- 4Ghz chip with mesh (logical and physical) 10x8 core FP, 1,28 TFlops
- Tile:
  - router: addesses each core on chip, implements the mesh
- VLIW processor (96 bit x instruction, up to 8 ops per cycle), in-order-execution, 32 registers (6Read/4Write), 2K Data, 3K Instruction cache, 2 FPU (9 stages, 2FLOPs/cycle sustained),
  - Cicli: FPU:9, Ld/St:2, Snd/ Rcv:2, Jmp/Br:1

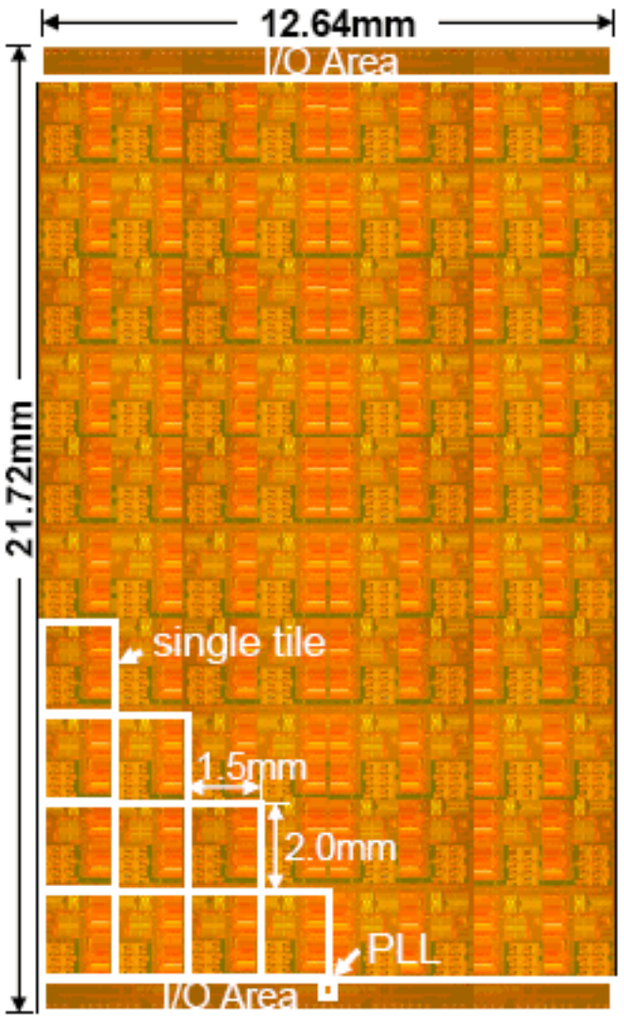

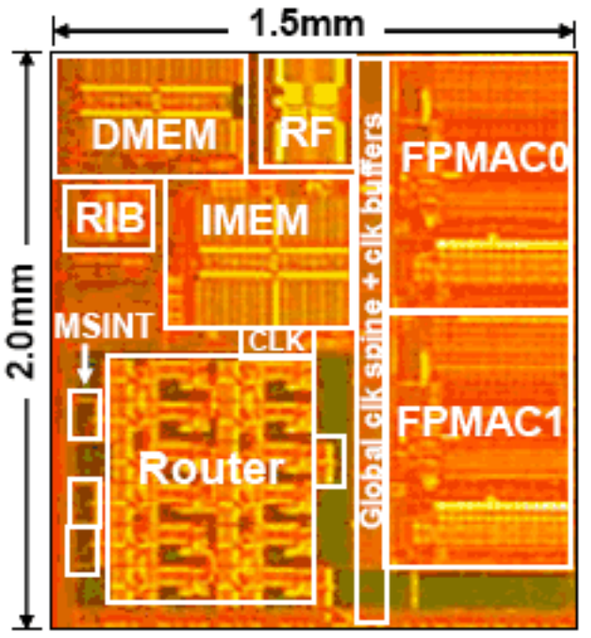

Technology	65nm CMOS Process
Interconnect	1 poly, 8 metal (Cu)
Transistors	100 Million
Die Area	275mm²
Tile area	3mm²
Package	1248 pin LGA, 14 layers, 343 signal pins

Intel trends

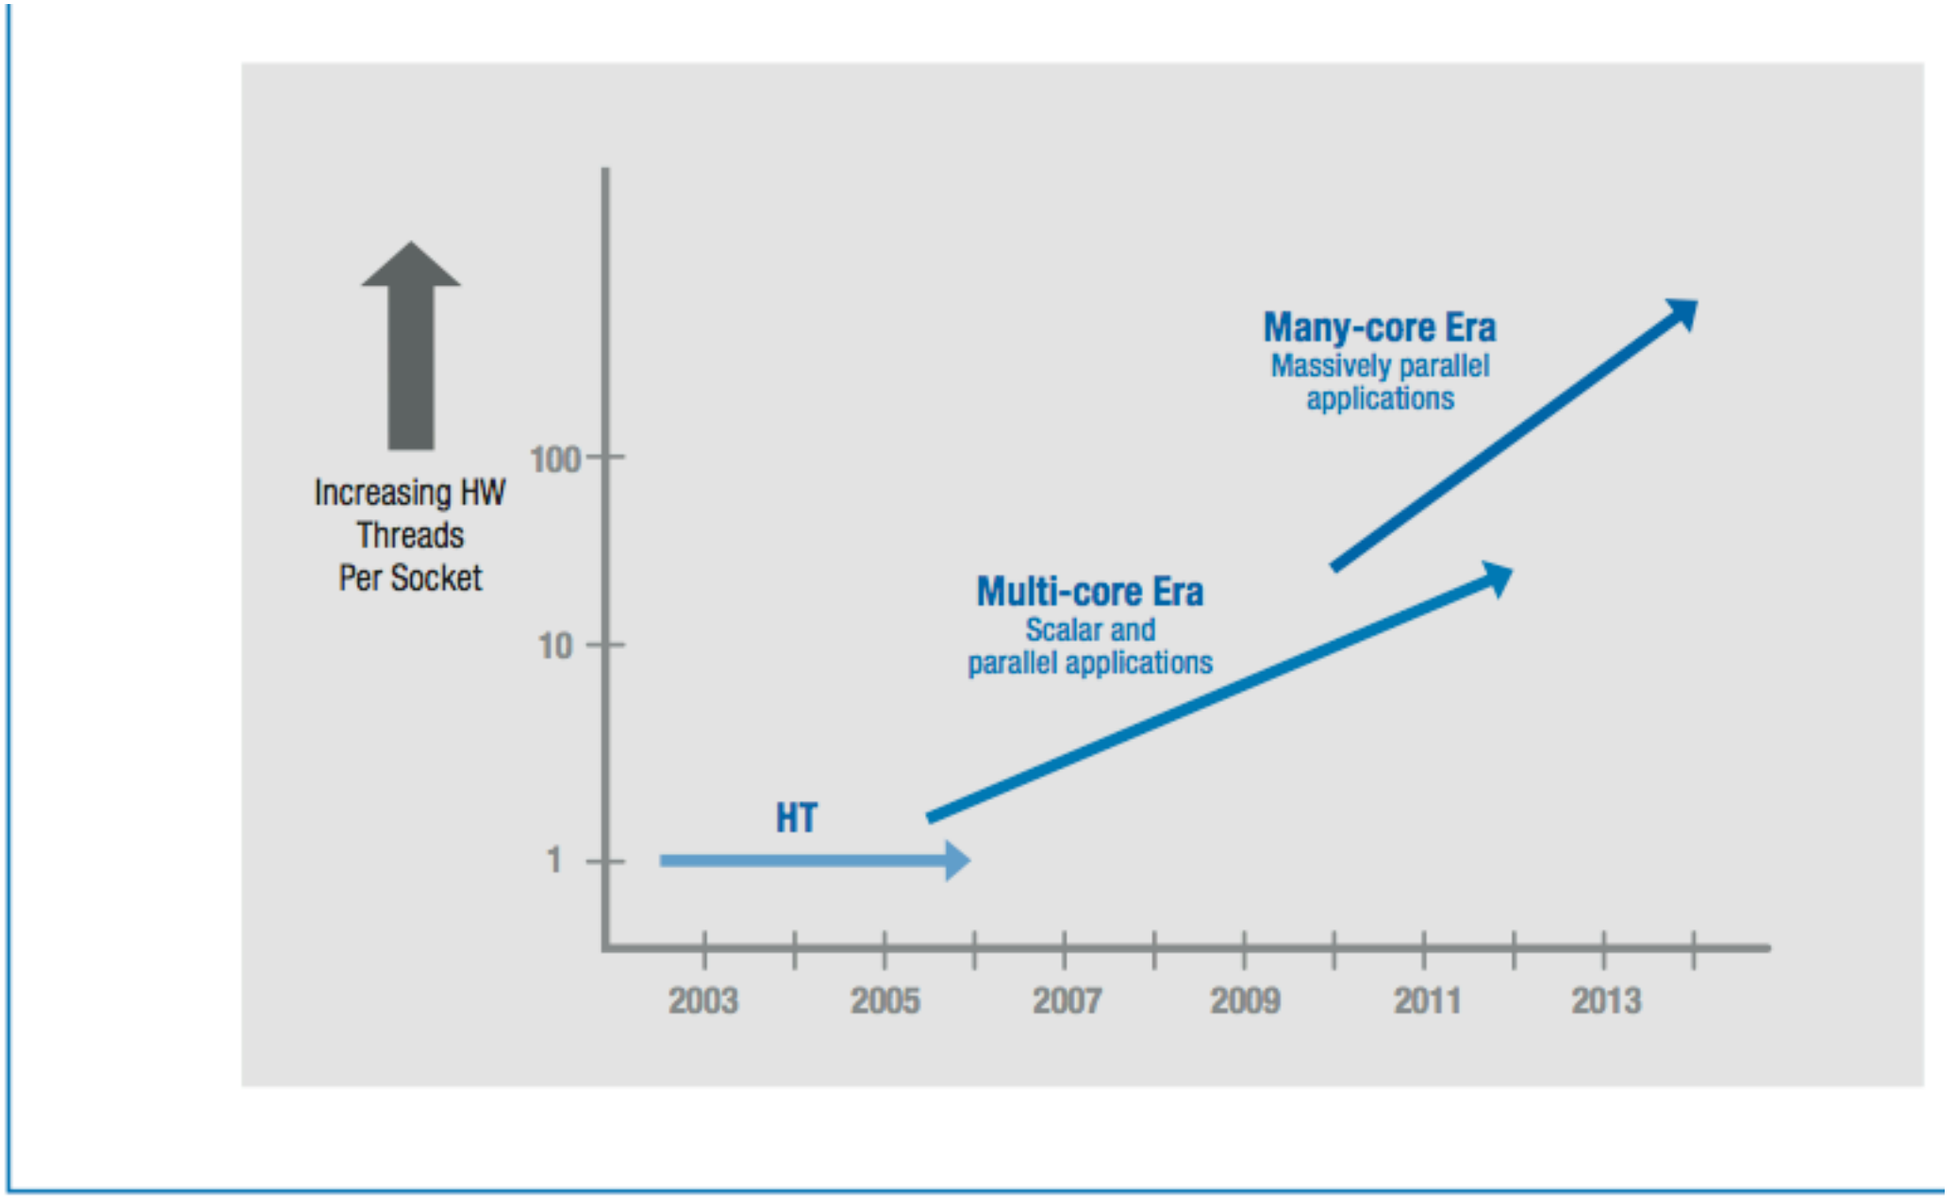

Figure 1: Current and expected eras of Intel® processor architectures

Intel trends (2)

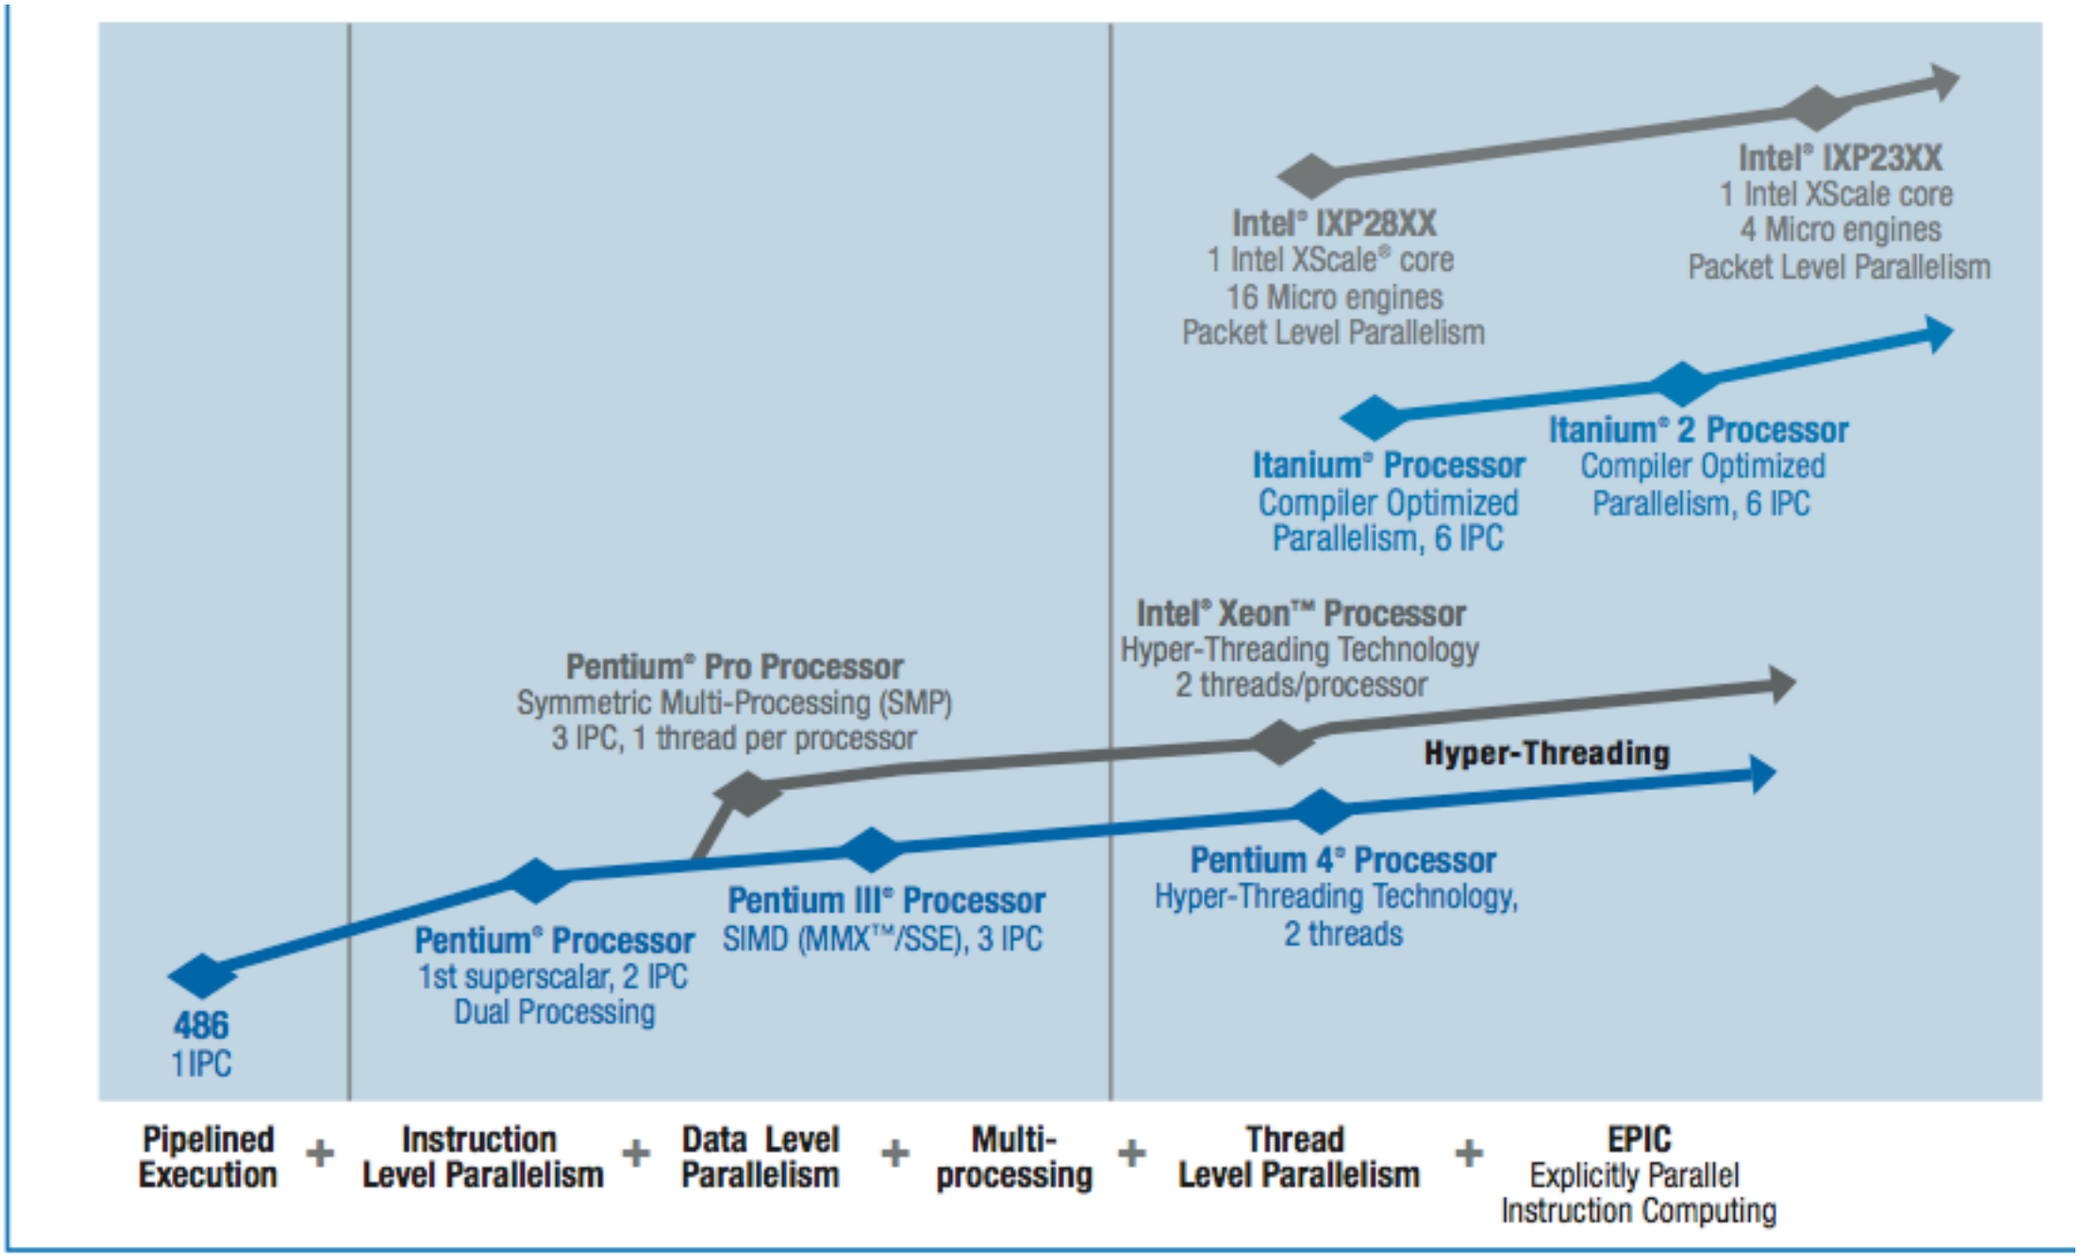

Figure 2: Driving increasing degrees of parallelism on Intel® processor architectures

Logistics: Earth Simulator (top500 1st 2 years ago)

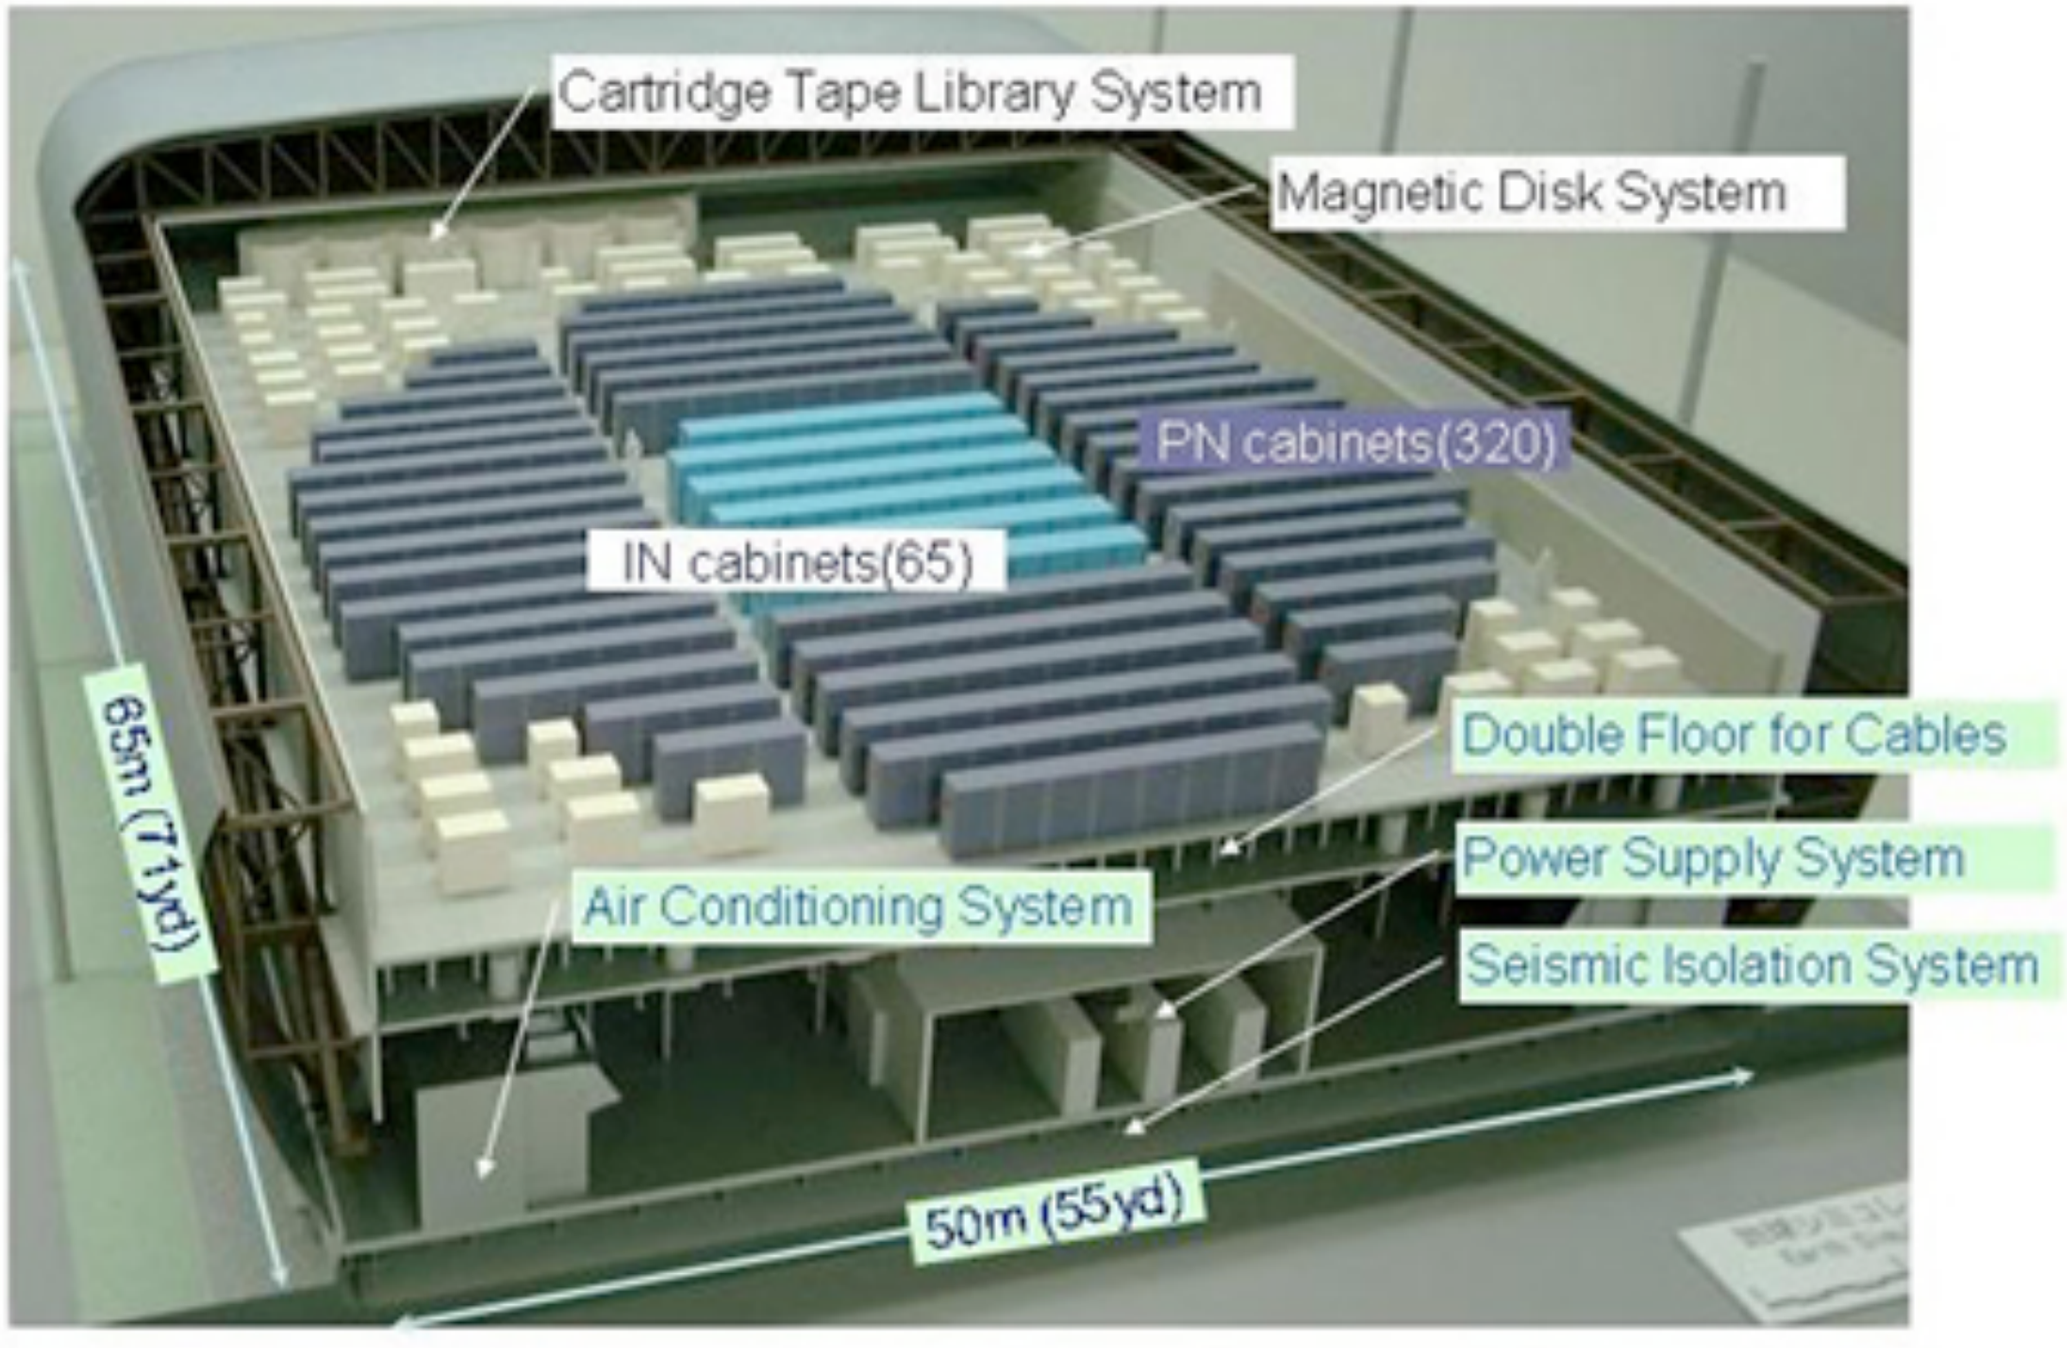

Green computing

Green computing

From Wikipedia, the free encyclopedia

Green computing is the study and practice of using computing resources efficiently. The primary objective of such a program is to account for the triple bottom line, an expanded spectrum of values and criteria for measuring organizational (and societal) success. The goals are similar to green chemistry; reduce the use of hazardous materials, maximize energy efficiency during the product's lifetime, and promote recyclability or biodegradability of defunct products and factory waste.

Data centers

Thinking Outside The Box, In The Box! Microsoft Begins Container Testing At Chicago Data Center

Microsoft's first Data Center that makes use of shipping containers as a primary server-packaging unit, has moved into the testing phase. The facility in Chicago is the first purpose-built Data Center to accommodate containers on such a large scale. When completed the Chicago Data Center will be one of the largest in the world with it's capacity to house hundreds of thousands of servers, each shipping container holds approx. 2,500 servers!

The Chicago Facility is much more energy efficient and environmentally friendly than Data Centers were before. All because the use of those 40ft containers! And energy efficiency isn't the only benefit, more about that in Mike Manos' point of view on the Chicago Data Center. Mike Manos is the general manager of data center services for Global Foundation Services/Microsoft.

> "..in my mind, containers are driving huge cost and efficiency (read also as cost benefits in addition to "green" benefits) gains for the business. This is an extremely important point, as Microsoft expands its data center infrastructure, it is supremely important that we follow an established smart growth methodology for our facilities that is designed to prevent overbuilding—and thus avoid associated costs to the environment and to our shareholders. We are a business after all. We must do all of this while also meeting the rapidly growing demand for Microsoft's Online and Live services."
>
> "Containers, and this new approach is definitely a change in how facilities have traditionally been developed, and as a result many people in our industry are intimidated by it. But they shouldn't be. Data center's have not changed in fundamental design for decades. Sometimes change is good. The exposure to any new idea is always met with resistance, but with a little education things change over time."

With that in mind Microsoft also announces they will hold their second Microsoft Data Center Experience (MDX) event in Chicago in Spring/summer 2009. The first one held in San Antonio was an opportunity for a couple hundred Microsoft enterprise customers to take a tour at the facilities and ask questions. That event was very successful and provided incredible feedback, so they decided to do one for the very first container data center too! More info on that on your Microsoft account management team or your local Microsoft sales office next Spring.

The newest Microsoft Data Center in West Des Moines, Iowa, will also be using this container technology.

Marketing ...

article | discussion | edit this page | history

MapReduce

From Wikipedia, the free encyclopedia

MapReduce is a software framework introduced by Google to support distributed computing on large data sets on clusters of computers [1]. The framework is inspired by map and reduce functions commonly used in functional programming,[2] although their purpose in the MapReduce framework is not the same as their original forms[3]. MapReduce libraries have been written in C++, Java, Python and other programming languages.

Overview [edit]

MapReduce is a framework for computing certain kinds of distributable problems using a large number of computers (nodes), collectively referred to as a cluster.

"Map" step: The master node takes the input, chops it up into smaller sub-problems, and distributes those to worker nodes. (A worker node may do this again in turn, leading to a multi-level tree structure.)

The worker node processes that smaller problem, and passes the answer back to its master node.

"Reduce" step: The master node then takes the answers to all the sub-problems and combines them in a way to get the output - the answer to the problem it was originally trying to solve.

The advantage of MapReduce is that it allows for distributed processing of the map and reduction operations. Provided each mapping operation is independent of the other, all maps can be performed in parallel - though in practice it is limited by the data source and/or the number of CPUs near that data. Similarly, a set of 'reducers' can perform the reduction phase - all that is required is that all outputs of the map operation which share the same key are presented to the same reducer, at the same time. While this process can often appear inefficient compared to algorithms that are more sequential, MapReduce can be applied to significantly larger datasets than that which "commodity" servers can handle - a large server farm can use MapReduce to sort a petabyte of data in only a few hours. The parallelism also offers some possibility of recovering from partial failure of servers or storage during the operation: if one mapper or reducer fails, the work can be rescheduled -assuming the input data is still available.

Marketing ...

article | discussion | edit this page | history

MapReduce

From Wikipedia, the free encyclopedia

MapReduce is a software framework introduced by Google to support distributed computing on large data sets on clusters of computers [1]. The framework is inspired by map and reduce functions commonly used in functional programming,[2] although their purpose in the MapReduce framework is not the same as their original forms[3]. MapReduce libraries have been written in C++, Java, Python and other programming languages.

Overview [edit]

MapReduce is a framework for computing certain kinds of distributable problems using a large number of computers (nodes), collectively referred to as a cluster.

"Map" step: The master node takes the input, chops it up into smaller sub-problems, and distributes those to worker nodes. (A worker node may do this again in turn, leading to a multi-level tree structure.)

The worker node pr

"Reduce" step: The
was originally trying

The advantage of M
independent of the
Similarly, a set of '
presented to the sa
MapReduce can be
a petabyte of data
if one mapper or re

Example

The canonical example application of MapReduce is a process to count the appearances of each different word in a set of documents:

```
map(String name, String document):
 // key: document name
 // value: document contents
 for each word w in document:
   EmitIntermediate(w, 1);

reduce(String word, Iterator partialCounts):
 // key: a word
 // values: a list of aggregated partial counts
 int result = 0;
 for each v in partialCounts:
   result += ParseInt(v);
 Emit(result);
```

Implementations

- The Google MapReduce framework is implemented in C++ with interfaces in Python and Java.
- Greenplum is a commercial MapReduce implementation, with support for Python, Perl, SQL and other languages.[7]
- GridGain is a free open source Java MapReduce implementation.
- The Hadoop project is a free open source Java MapReduce implementation.
- Phoenix [1] is a shared-memory implementation of MapReduce implemented in C.
- FileMap is an open version of the framework that operates on files using existing file-processing tools rather than tuples.
- MapReduce has also been implemented for the Cell Broadband Engine, also in C. [2]
- MapReduce has been implemented on NVIDIA GPUs (Graphics Processors) using CUDA [3].
- Qt Concurrent is a simplified version of the framework, implemented in C++, used for distributing a task between multiple processor cores.
- CouchDB uses a MapReduce framework for defining views over distributed documents
- Skynet is an open source Ruby implementation of Google's MapReduce framework
- Disco is an open source MapReduce implementation by Nokia. Its core is written in Erlang and jobs are normally written in Python.
- Aster Data Systems *n*Cluster In-Database MapReduce implements MapReduce inside the database.
- The open-source Hive framework from Facebook (which provides a SQL-like language over files, layered on the open-source Hadoop MapReduce engine.)

cluster.

"Map" step: The master node takes the input, chops it up into smaller sub-problems, and distributes those to worker nodes. (A worker node may do this again in turn, leading to a multi-level tree structure.)

The worker node pr

"Reduce" step: The was originally trying

The advantage of M independent of the Similarly, a set of ' presented to the sa MapReduce can be a petabyte of data if one mapper or re

Example

The canonical example application of MapReduce is a process to count the appearances of each different word in a set of documents:

```
map(String name, String document):
  // key: document name
  // value: document contents
  for each word w in document:
    EmitIntermediate(w, 1);

reduce(String word, Iterator partialCounts):
  // key: a word
  // values: a list of aggregated partial counts
  int result = 0;
  for each v in partialCounts:
    result += ParseInt(v);
  Emit(result);
```

MapReduce Execution Overview

The Map invocations are distributed across multiple machines by automatically partitioning the input data into a set of M splits or *shards*. The input shards can be processed in parallel on different machines.

Reduce invocations are distributed by partitioning the intermediate key space into R pieces using a partitioning function (e.g., hash(key) mod R). The number of partitions (R) and the partitioning function are specifed by the user.

The illustration below shows the overall fow of a MapReduce operation. When the user program calls the MapReduce function, the following sequence of actions occurs (the numbered labels in the illustration correspond to the numbers in the list below).

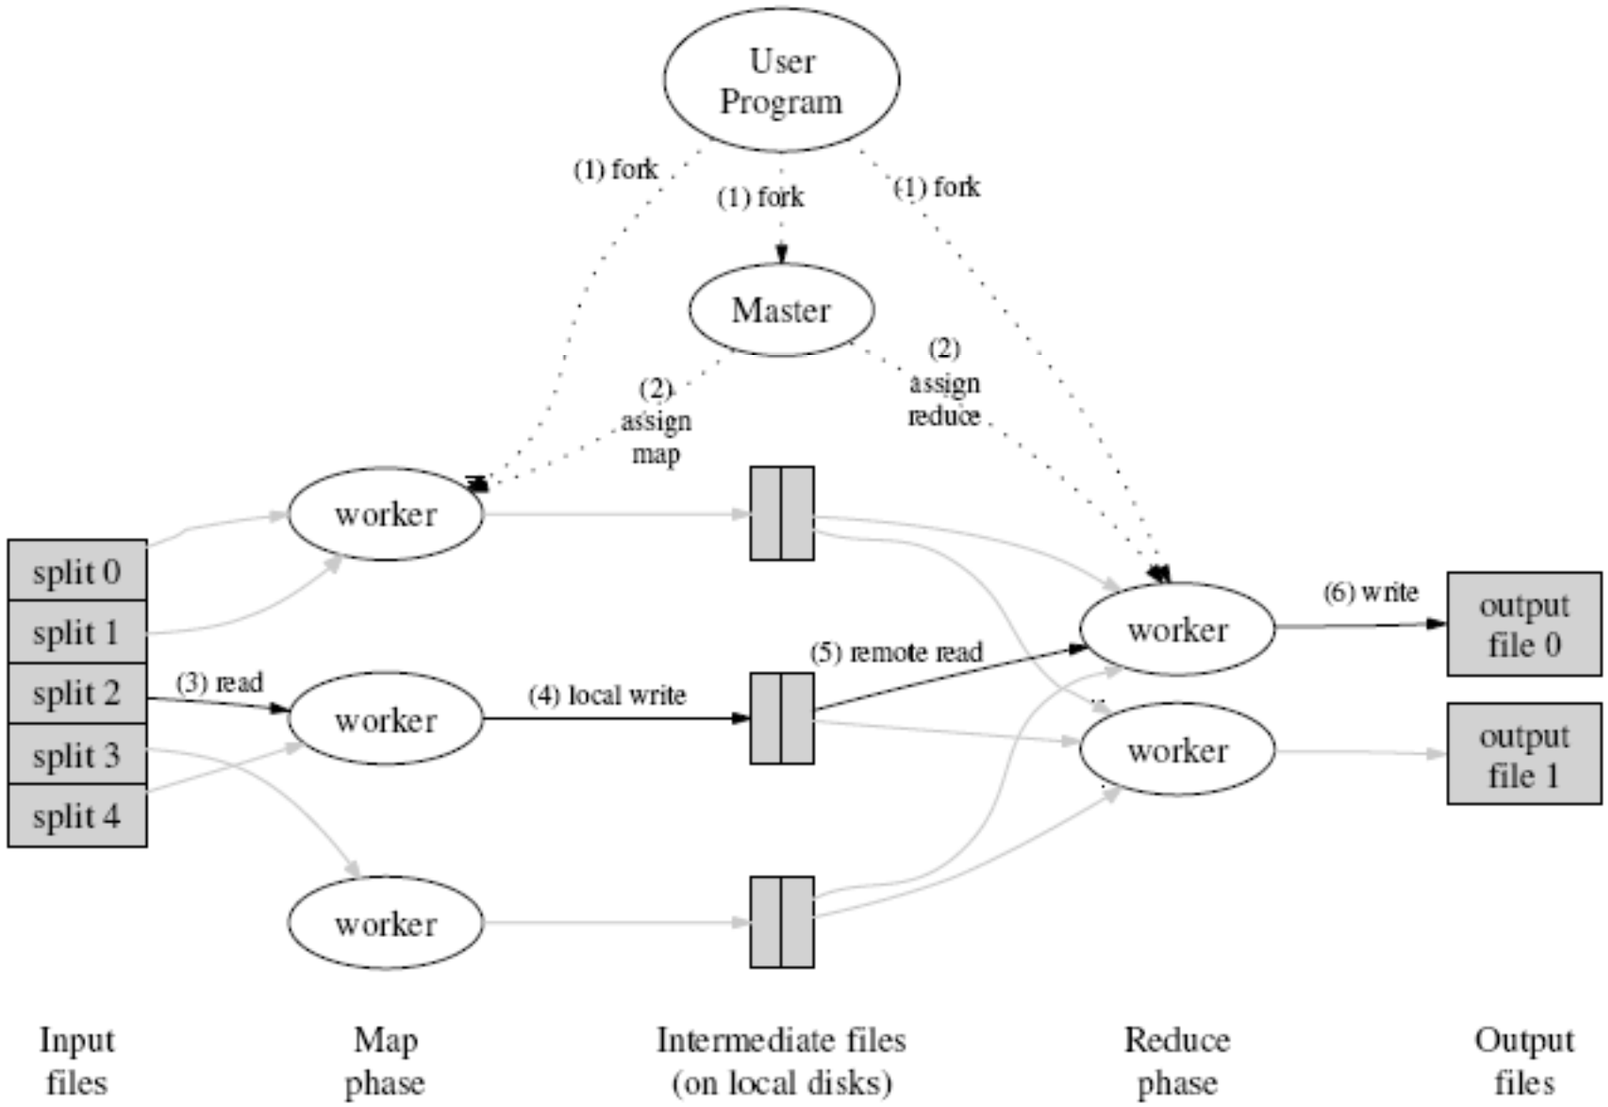

Apache

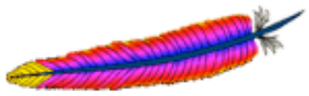


hadoop

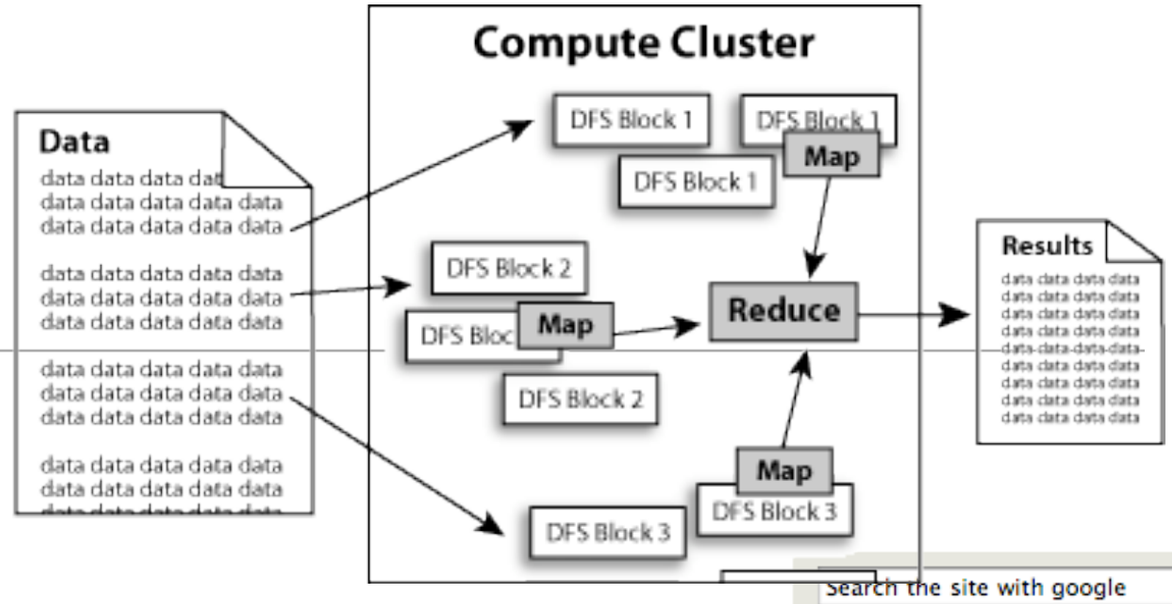
Compute Cluster
Data
data data data data
data data data data data
data data data data data
data data data data data
data data data data data
data data data data data
data data data data data
data data data data data
DFS Block 1
DFS Block 1
DFS Block 1
Map
DFS Block 2
DFS Bloc
Map
Reduce
DFS Block 2
Map
DFS Block 3
DFS Block 3
Results
data data data data
data data data data
data data data data
data data data data
data-data-data-data-
data data data data
data data data data
data data data data
data data data data

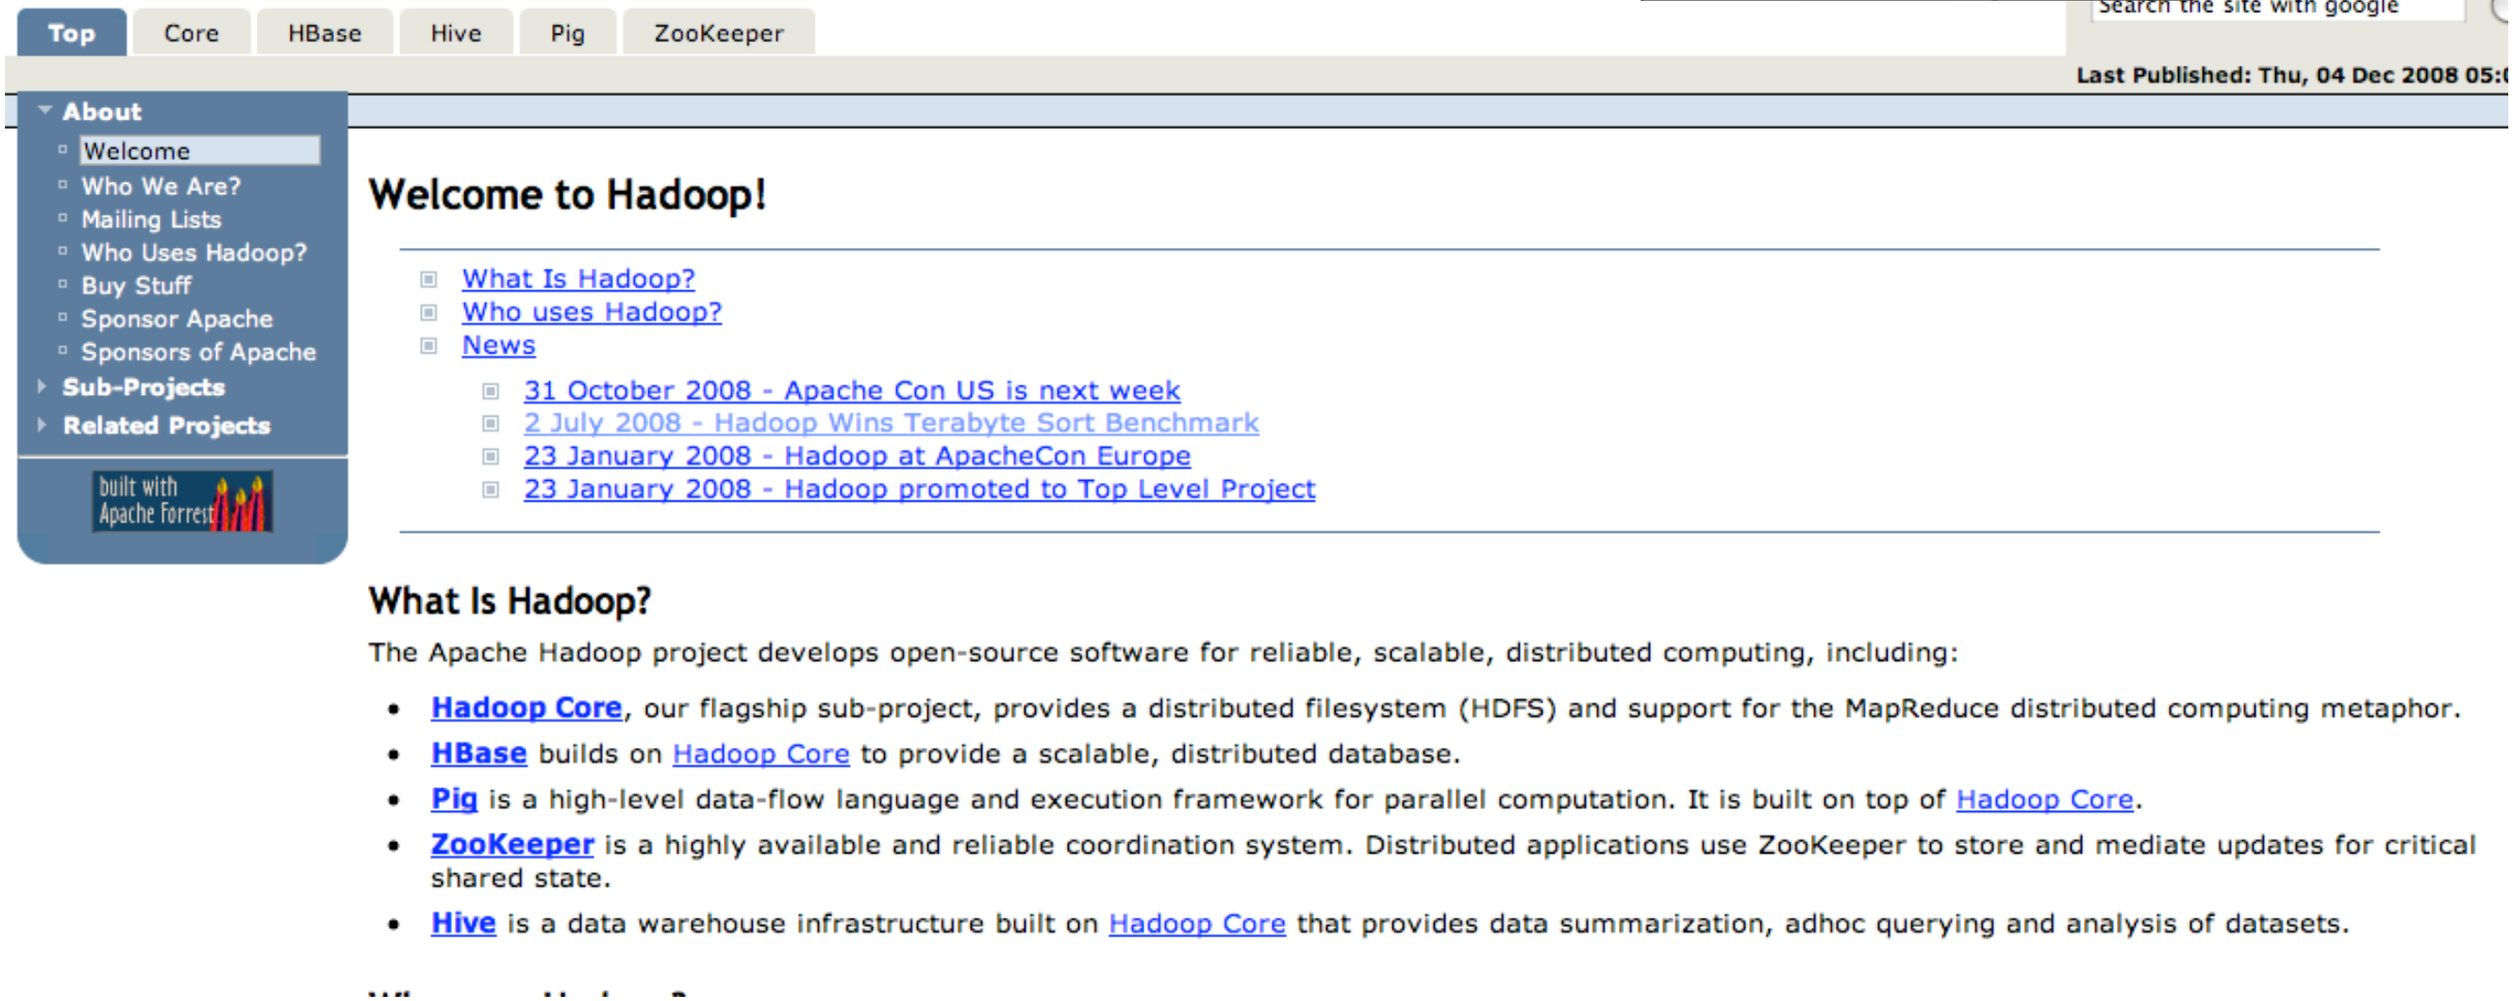
Top
Core
HBase
Hive
Pig
ZooKeeper
Search the site with google
Last Published: Thu, 04 Dec 2008 05:
About
Welcome
Who We Are?
Mailing Lists
Who Uses Hadoop?
Buy Stuff
Sponsor Apache
Sponsors of Apache
Sub-Projects
Related Projects
built with Apache Forrest
Welcome to Hadoop!
What Is Hadoop?
Who uses Hadoop?
News
31 October 2008 - Apache Con US is next week
2 July 2008 - Hadoop Wins Terabyte Sort Benchmark
23 January 2008 - Hadoop at ApacheCon Europe
23 January 2008 - Hadoop promoted to Top Level Project
What Is Hadoop?
The Apache Hadoop project develops open-source software for reliable, scalable, distributed computing, including:
Hadoop Core, our flagship sub-project, provides a distributed filesystem (HDFS) and support for the MapReduce distributed computing metaphor.
HBase builds on Hadoop Core to provide a scalable, distributed database.
Pig is a high-level data-flow language and execution framework for parallel computation. It is built on top of Hadoop Core.
ZooKeeper is a highly available and reliable coordination system. Distributed applications use ZooKeeper to store and mediate updates for critical shared state.
Hive is a data warehouse infrastructure built on Hadoop Core that provides data summarization, adhoc querying and analysis of datasets.

Cloud computing

Cloud computing

Grid Computing

Metacomputing

What is Metacomputing?

Metacomputing involves using multiple computers linked by high-speed networks to solve application problems; they appear to the user as a single computer but provide a level of performance far in excess of individual computers in the network. Metacomputers or computational grids function like a networked virtual supercomputer. Metacomputing was born out of a need to utilize greater processing power than was available in a single site, and to combine the power of computers with different architecture.

Cloud computing

Cloud computing

Grid computing (or the use of a computational grid) is the application of several computers to a single problem at the same time -- usually to a scientific or technical problem that requires a great number of computer processing cycles or access to large amounts of data. According to John Patrick, IBM's vice president for Internet strategies, "the next big thing will be grid computing."[*citation needed*]

Grid computing depends on software to divide and apportion pieces of a program among several computers, sometimes up to many thousands. Grid computing can also be thought of as distributed[*citation needed*] and large-scale cluster computing, as well as a form of network-distributed parallel processing[*citation needed*]. It can be small -- confined to a network of computer workstations within a corporation, for example -- or it can be a large, public collaboration across many companies or networks.

It is a form of distributed computing whereby a "super and virtual computer" is composed of a cluster of networked, loosely coupled computers, acting in concert to perform very large tasks. This technology has been applied to computationally intensive scientific, mathematical, and academic problems through volunteer computing, and it is used in commercial enterprises for such diverse applications as drug discovery, economic forecasting, seismic analysis, and back-office data processing in support of e-commerce and Web services.

What distinguishes grid computing from conventional cluster computing systems is that grids tend to be more loosely coupled, heterogeneous, and geographically dispersed. Also, while a computing grid may be dedicated to a specialized application, it is often constructed with the aid of general-purpose grid software libraries and middleware.

Metacomputing

What is Metacomputing?

Metacomputing involves using multiple computers linked by high-speed networks to solve application problems; they appear to the user as a single computer but provide a level of performance far in excess of individual computers in the network. Metacomputers or computational grids function like a networked virtual supercomputer. Metacomputing was born out of a need to utilize greater processing power than was available in a single site, and to combine the power of computers with different architecture.

Cloud c

Cloud computing refers to the use of Internet ("cloud") based computer technology for a variety of services.[1] It is a style of computing in which dynamically scalable and often virtualised resources are provided as a service over the Internet[2][3][4][5]. Users need not have knowledge of, expertise in, or control over the technology infrastructure "in the cloud" that supports them[6].

The concept incorporates software as a service (SaaS), Web 2.0 and other recent, well-known technology trends, in which the common theme is reliance on the Internet for satisfying the computing needs of the users. Often-quoted examples are Salesforce.com and Google Apps which provide common business applications online that are accessed from a web browser, while the software and data are stored on the servers.

The *cloud* is a metaphor for the Internet, based on how it is depicted in computer network diagrams, and is an abstraction for the complex infrastructure it conceals.[7]

Grid computing (or the use of a computational grid) is the application of several computers to a single problem at the same time -- usually to a scientific or technical problem that requires a great number of computer processing cycles or access to large amounts of data. According to John Patrick, IBM's vice president for Internet strategies, "the next big thing will be grid computing."[*citation needed*]

Grid computing depends on software to divide and apportion pieces of a program among several computers, sometimes up to many thousands. Grid computing can also be thought of as distributed[*citation needed*] and large-scale cluster computing, as well as a form of network-distributed parallel processing[*citation needed*]. It can be small -- confined to a network of computer workstations within a corporation, for example -- or it can be a large, public collaboration across many companies or networks.

It is a form of distributed computing whereby a "super and virtual computer" is composed of a cluster of networked, loosely coupled computers, acting in concert to perform very large tasks. This technology has been applied to computationally intensive scientific, mathematical, and academic problems through volunteer computing, and it is used in commercial enterprises for such diverse applications as drug discovery, economic forecasting, seismic analysis, and back-office data processing in support of e-commerce and Web services.

What distinguishes grid computing from conventional cluster computing systems is that grids tend to be more loosely coupled, heterogeneous, and geographically dispersed. Also, while a computing grid may be dedicated to a specialized application, it is often constructed with the aid of general-purpose grid software libraries and middleware.

Metacomputing

What is Metacomputing?

Metacomputing involves using multiple computers linked by high-speed networks to solve application problems; they appear to the user as a single computer but provide a level of performance far in excess of individual computers in the network. Metacomputers or computational grids function like a networked virtual supercomputer. Metacomputing was born out of a need to utilize greater processing power than was available in a single site, and to combine the power of computers with different architecture.