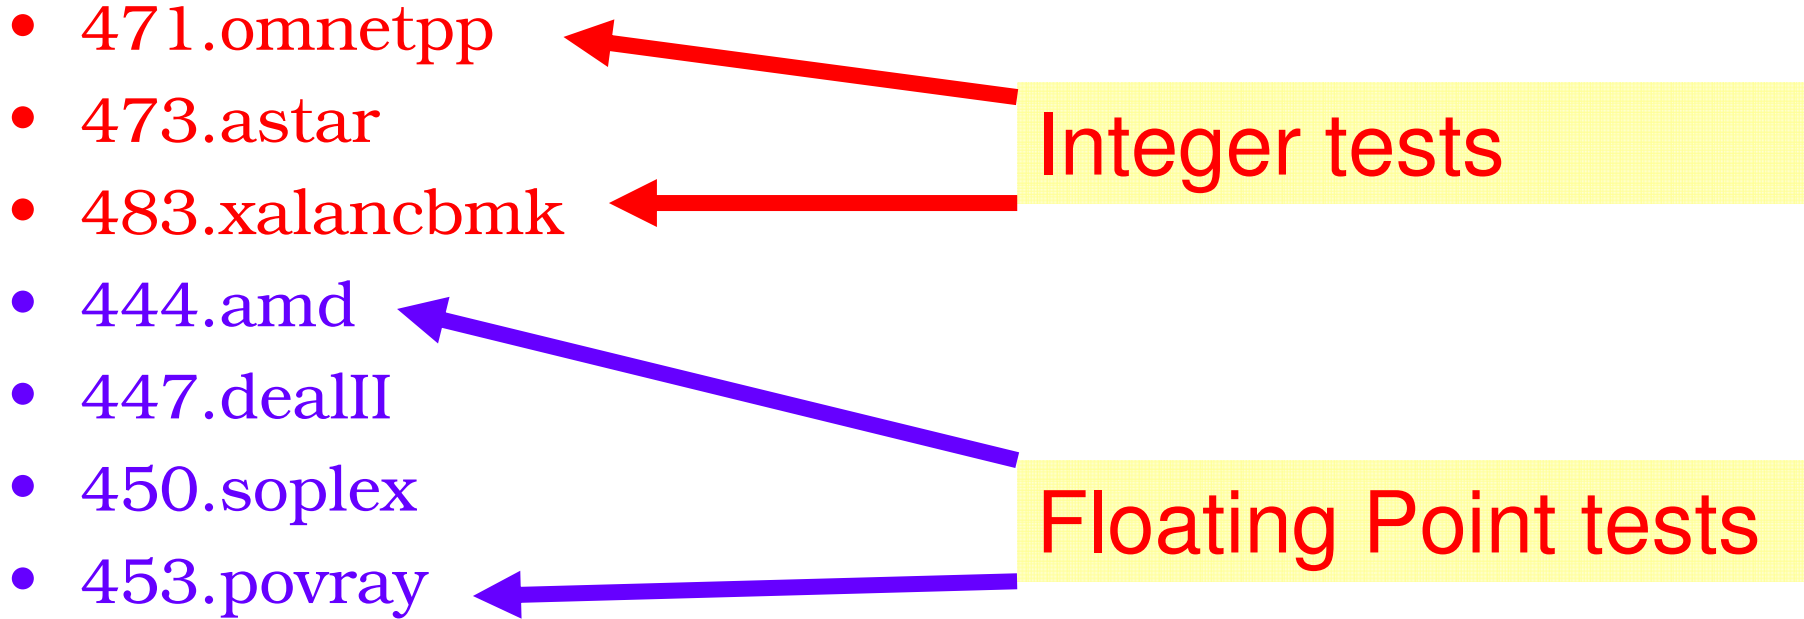


INFN - Padova

michele.michelotto at pd.infn.it

HS06 on last generation of HEP  
worker nodes


Berkeley, Hepix Fall '09



- Geometric average on 7 tests. Sum on all cores
- Sum on all Worker nodes

Worker Node, cpu, clock, L2 cache, - Main Memory	SPECint2000 gcc	SPEC CPU2006 Int - gcc	HEP-SPEC-06	site measured
2 x Nocona/Irvindale 2.8 GHz/1 MB - 2GB	1501	11.06	10.24	CERN
2 x Nocona/Irvindale 2.8 GHz/2 MB - 4GB	1495	10.09	9.63	CERN
2 x Nocona/Irvindale 2.8 GHz/2 MB - 2GB	1673	11.87		CERN
2 x Nocona/Irvindale 2.8 GHz/2 MB - 2GB	1703	12.26		CERN
2 X Opteron 275 2.2 GHz/2 MB - 4GB	4133	28.76	28.03	CERN
2 x Woodcrest 2.66 GHz/4 MB - 8GB	5675	36.77	35.58	CERN
2 x Woodcrest 3.00 GHz/4 MB - 8GB	6181	39.39	38.21	CERN
2 x Opteron 2218 (Rev. F) 2.6 GHz/2 MB- 8GB	4569	31.4	31.67	CERN
2 x Clovertown 2.33 GHz/2x4MB - 16GB	9462	60.89	57.52	CERN
2 x Harpertown (E5410) 2.33 GHz/2x6 MB - 16GB	10556	64.78	60.76	CERN
2 x Harpertown (5440) 2.83 GHz/2x6MB - 16GB	11850	73.32		DESY
2 x Harpertown (5410) 2.33 GHz/2x6MB - 16GB	11164	65.93	62.12	INFN-PD
2 x Barcelona (2352) 2.10 GHz/2x4MB - 16GB	8488	56.23		INFN-PD
2 x Barcelona (2360) 2.50 GHz/2x4MB - 16GB	9939	63.75	63.19	INFN-PD
2 x Barcelona (2356) 2.30 GHz/4x512 KB - 16GB	9565	61.05	59.74	GridKa
2 x Shanghai (2376) 2.30 GHz/4x512KB+6MB L3 - 16GB	10962	66.88	65.85	GridKa
2 x Harpertown (E5430) 2.66 GHz/2x6MB -16GB	12122	72.14	68.04	GridKa
2 x Opteron (2216) 2.4GHz -8GB		21.86	22.22	RAL
2 x Barcelona (2354) 2.2 GHz -16GB		58.17	58.1	RAL
2 x Clovertown (E5335) 2.00 GHz/2x4 MB - 16GB		54.42	52.31	RAL
2 x Harpertown (L5410) 2.33 GHz/2x6MB -16GB		68.84	62.23	RAL
2 x Harpertown (E5410) 2.33 GHz/2x6MB -16GB		65.73	62.21	RAL
2 x Harpertown (L5420) 2.50 GHz/2x6MB - 16GB		68.73	62.11	RAL
2 x Harpertown (E5420) 2.50 GHz/2x6MB - 16GB		69.13	65.11	RAL
2 x Harpertown (E5420) 2.50 GHz/2x6MB, 16GB		66.83	64.85	RAL
2 x Harpertown (E5440) 2.83 GHz/2x6MB, 16GB		74.92	68.51	RAL

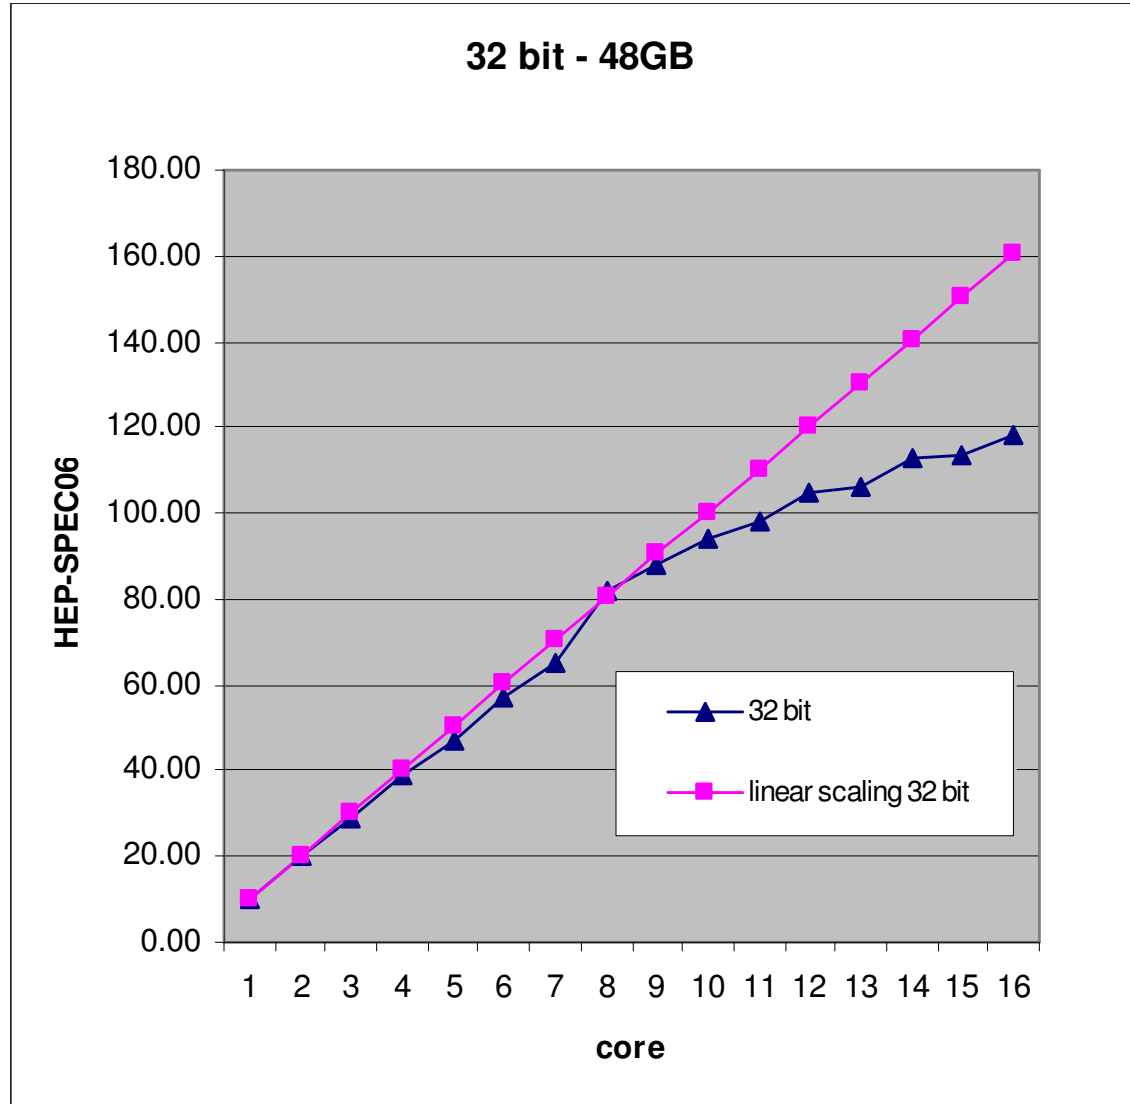
processor	HS06	HS06/Clock	HS06/Clock/ core	Core/ Logical cpu
Intel Clovertown 53xx	53-60	23-26	2.90-3.20	8
Intel Harpertown 54xx	60-70	25-28	3.20-3.50	8
AMD Shanghai 23xx	60-74	25-27	3.20-3.60	8
AMD Istanbul 24xx	96-99	40-44	3.34 -3.73	12
Intel Gainestown 5520	80-95-120	43-53	3.33-5.39	8-16



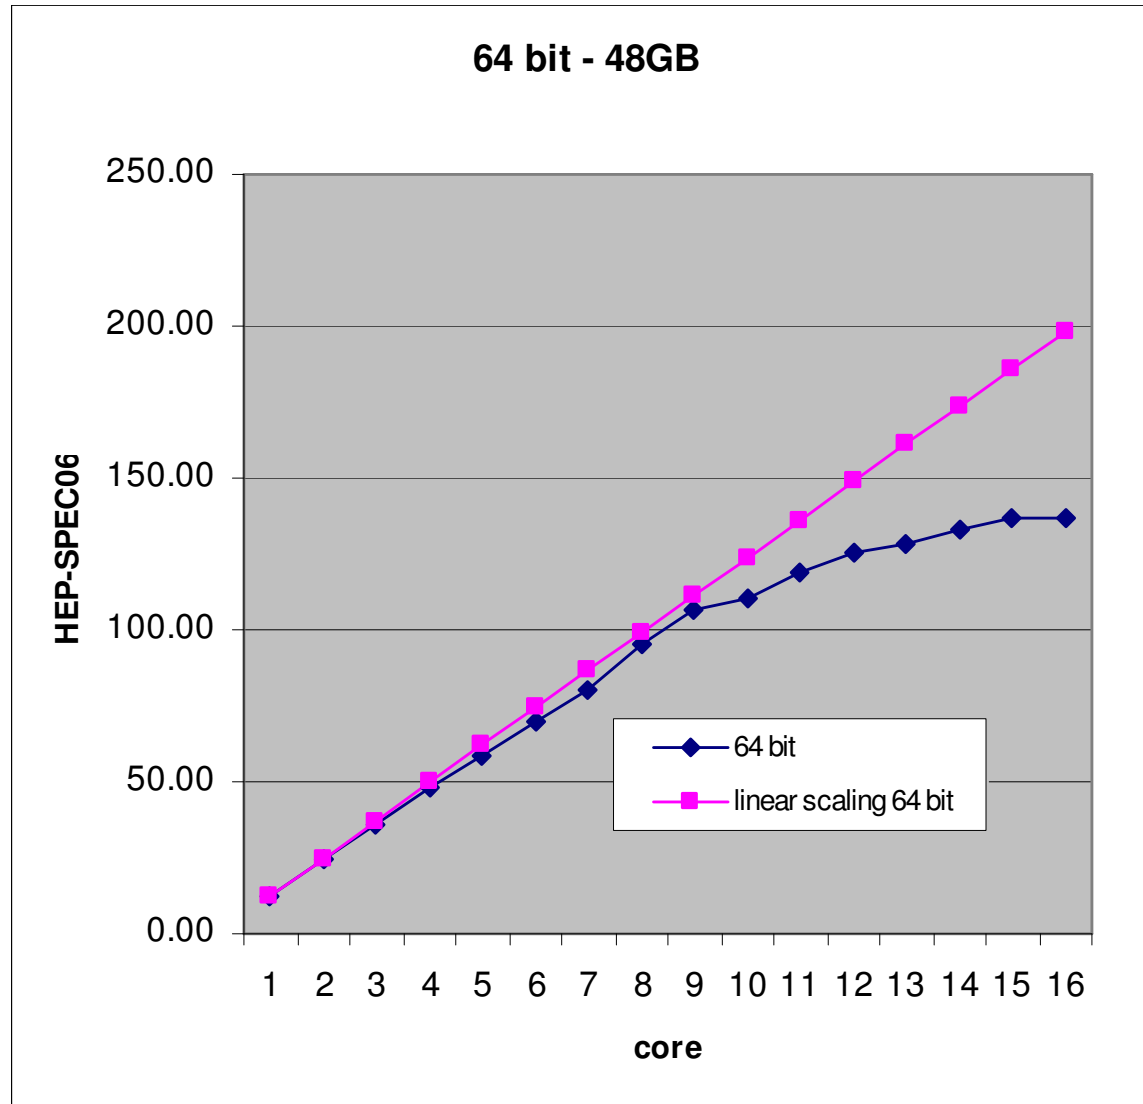
- About 60 Measurement from Brunengo, Macorini, Crescente, Calzolari (all INFN), Srinivasan (LBNL), Iribarren (CERN), Alef (GridKa), PIC, RAL, Nikhef

- Phase space complicated
  - 32-64 bit
  - Hyperthreading ON - OFF
  - Memory (Size and number of channels)
  - Turbo Mode ON - OFF

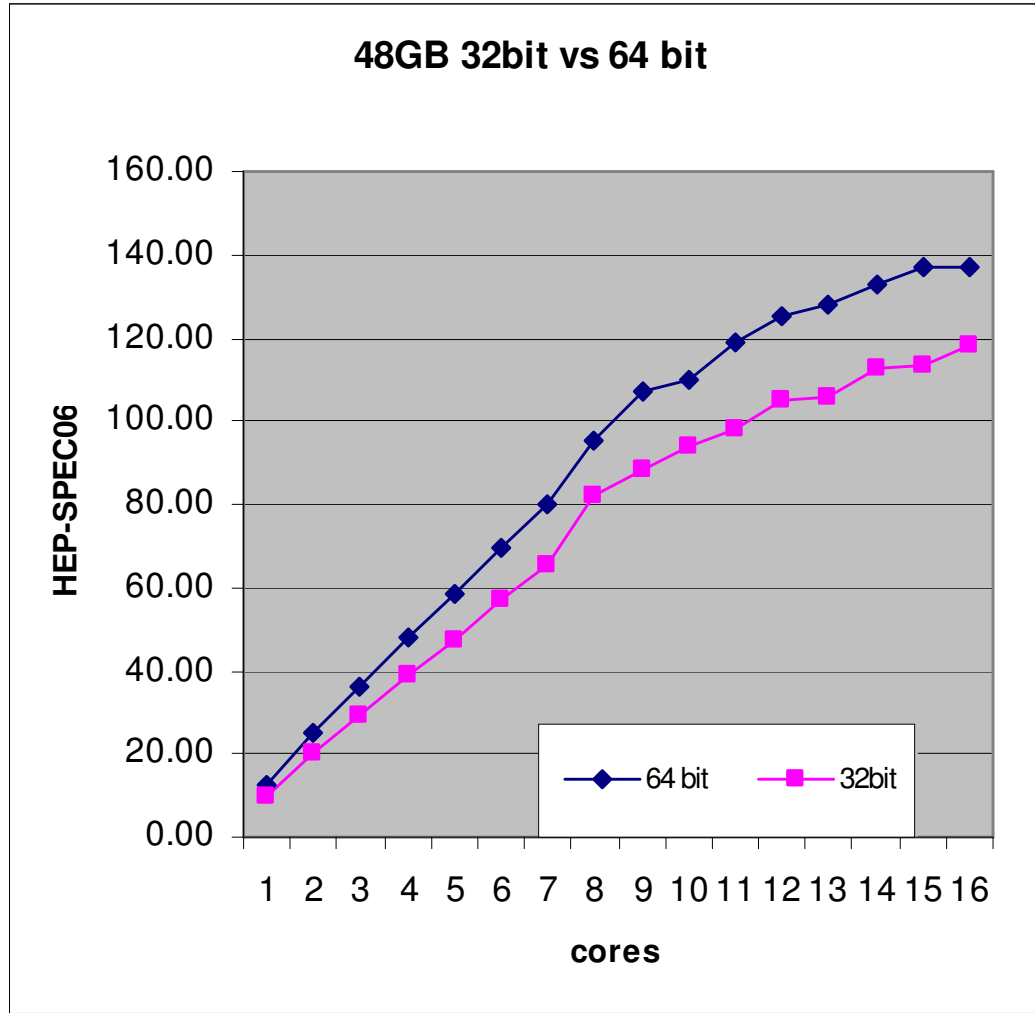
•HS06 HT  
is on  
118.30 (16t)  
81.81  
(8t)



•HS06 HT  
is on  
136.74  
(16t)  
95.14  
(8t)

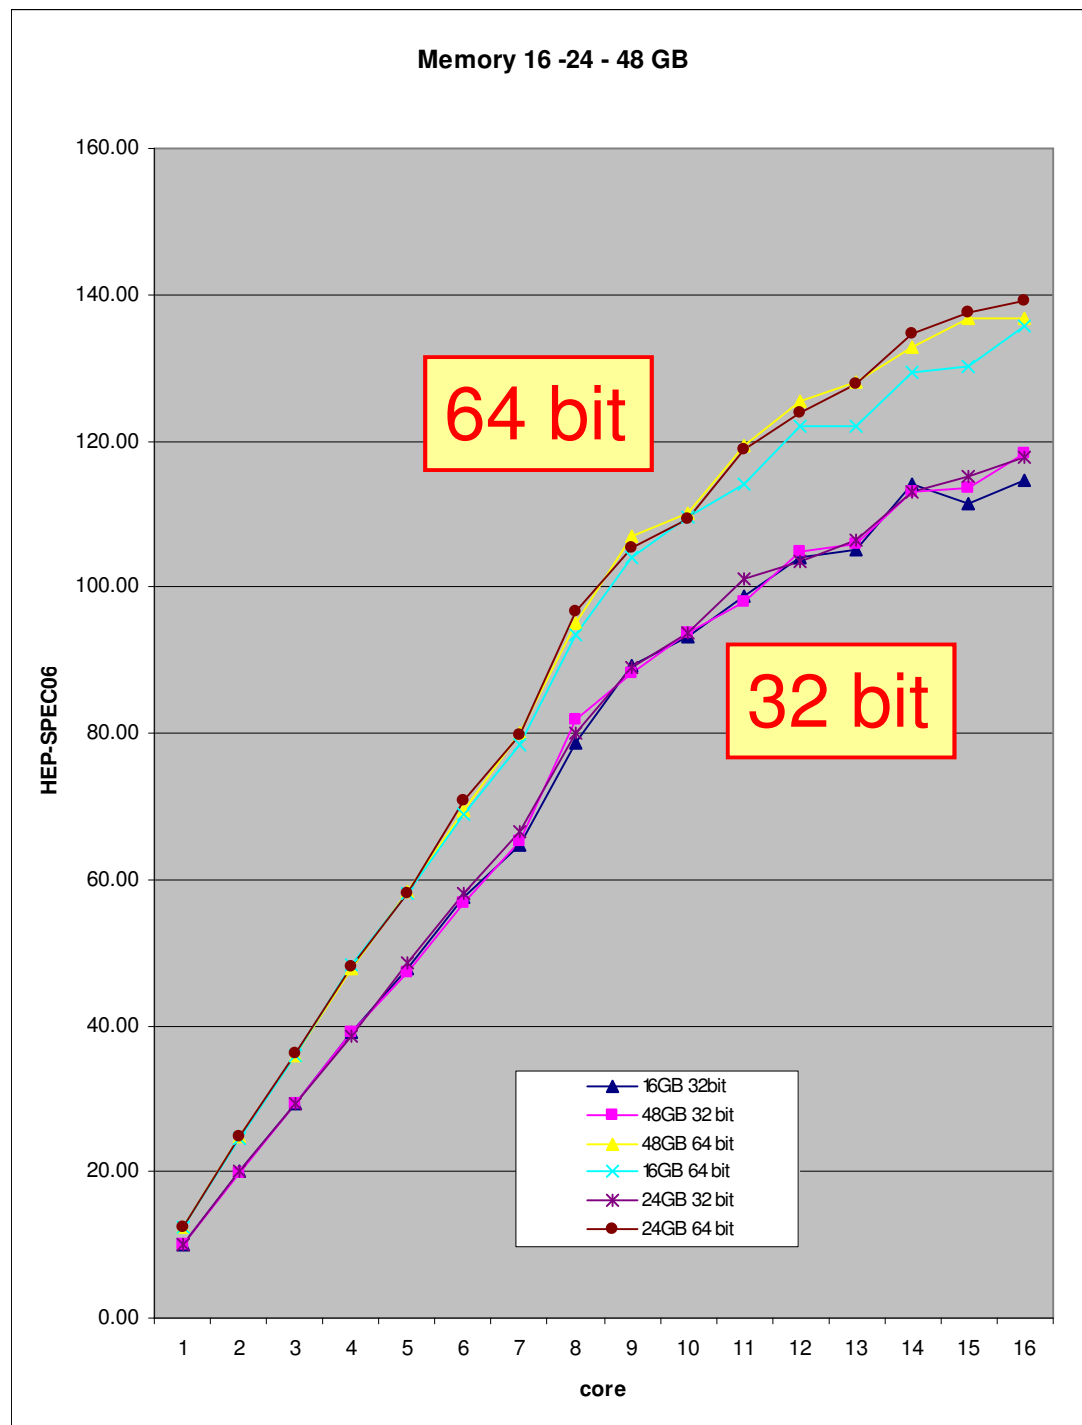


# 32 vs 64


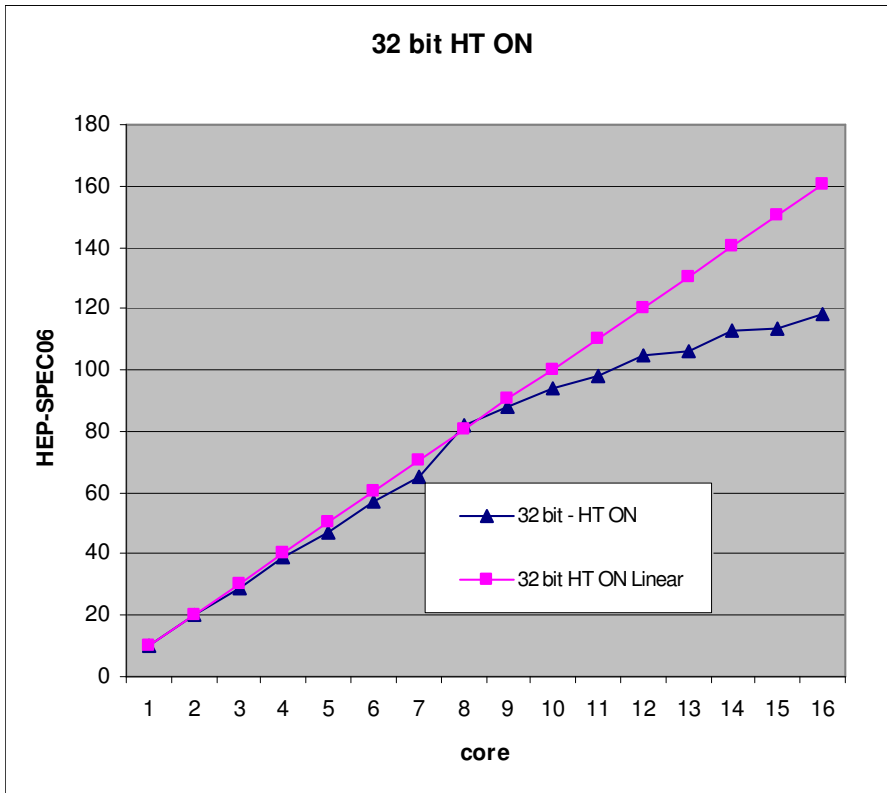





Memory  
 48GB (6x8)  
 24GB (3x8)  
 16GB (2x8)




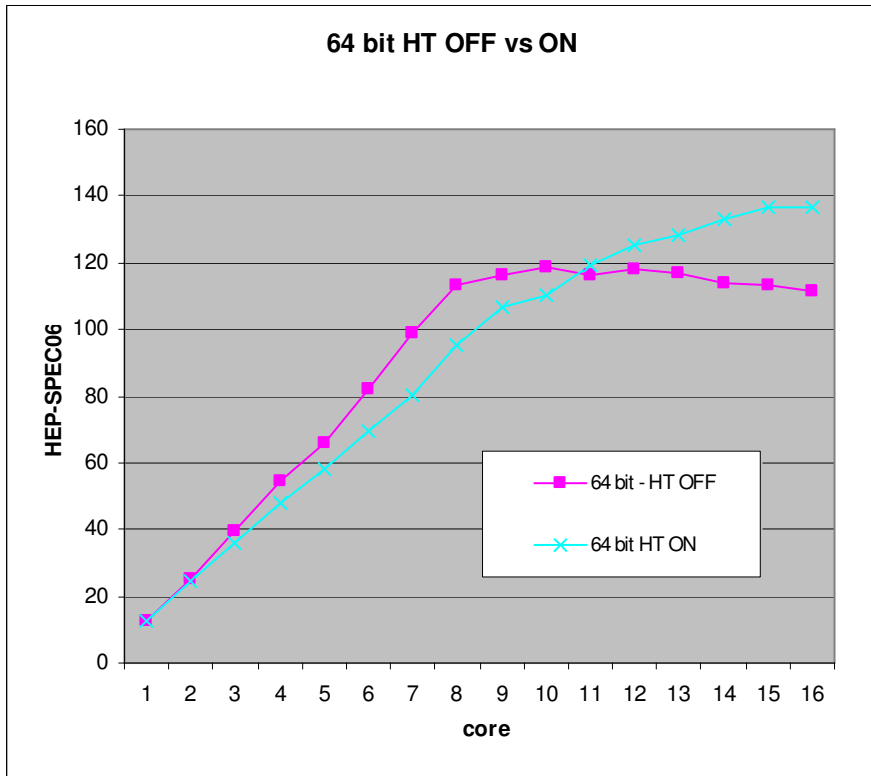
- With HT OFF I'd imagine saturation after 8<sup>th</sup> thread
- Will HT OFF 1-8 = HT ON?






HT ON:81.81  
 HT OFF: 95.96  
 HT OFF is  
 better up to 11t

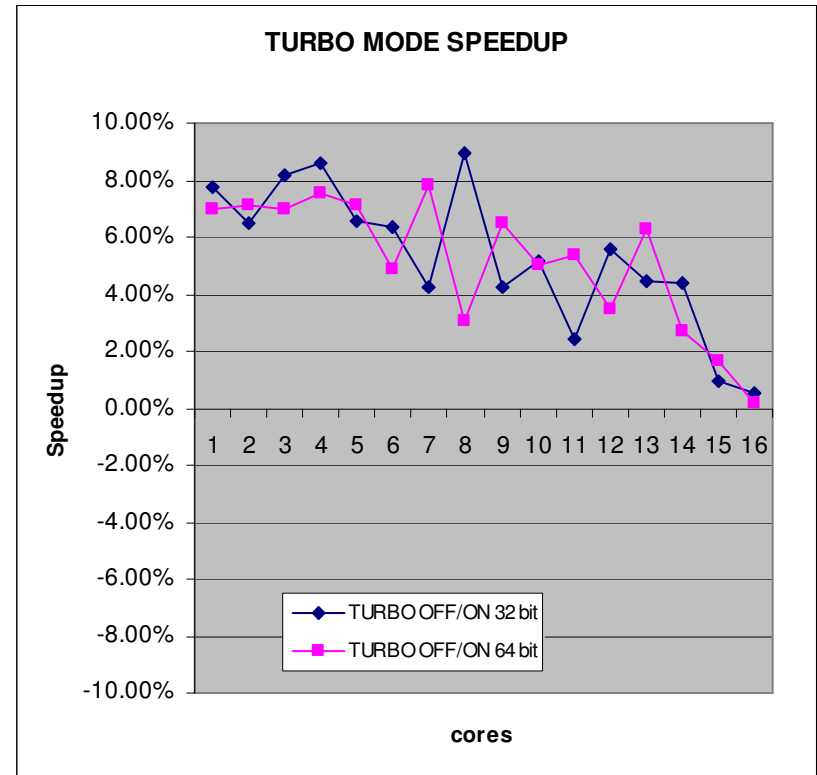

- At 64 bit: same behaviour




- Turn off the voltage on idle cores and overclock +133 MHz or even +266MHz the active cores if temperature is ok
  - 5520 Default clock is 2266 MHz
  - 5520 3 of 4 core +1bin → 2400 MHz
  - 5520 1 or 2 core +2bin → 2533 MHz
- New half-generation e.g. Xeon 3500 up to 4 bin



- Benefit of Turbo mode decrease when number of active core increase




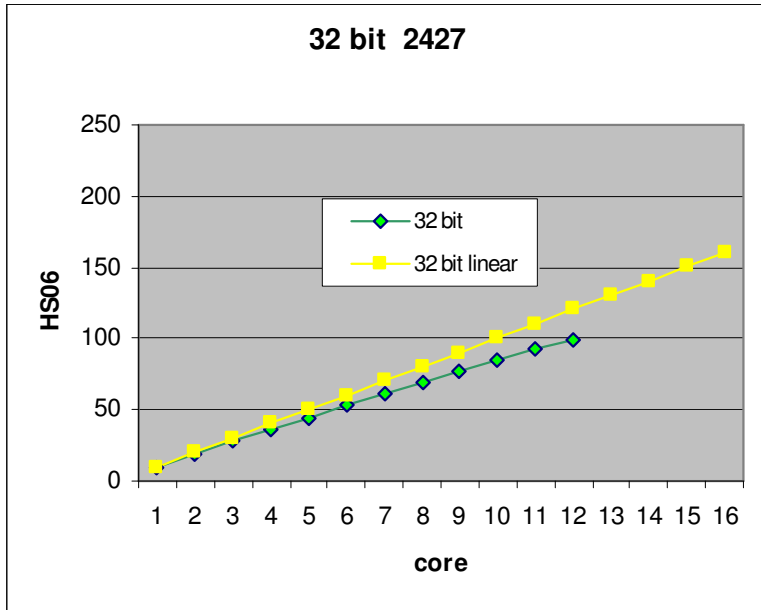
•HS06:  
98.05 (12t)



•HS06:  
111.46 (12t)



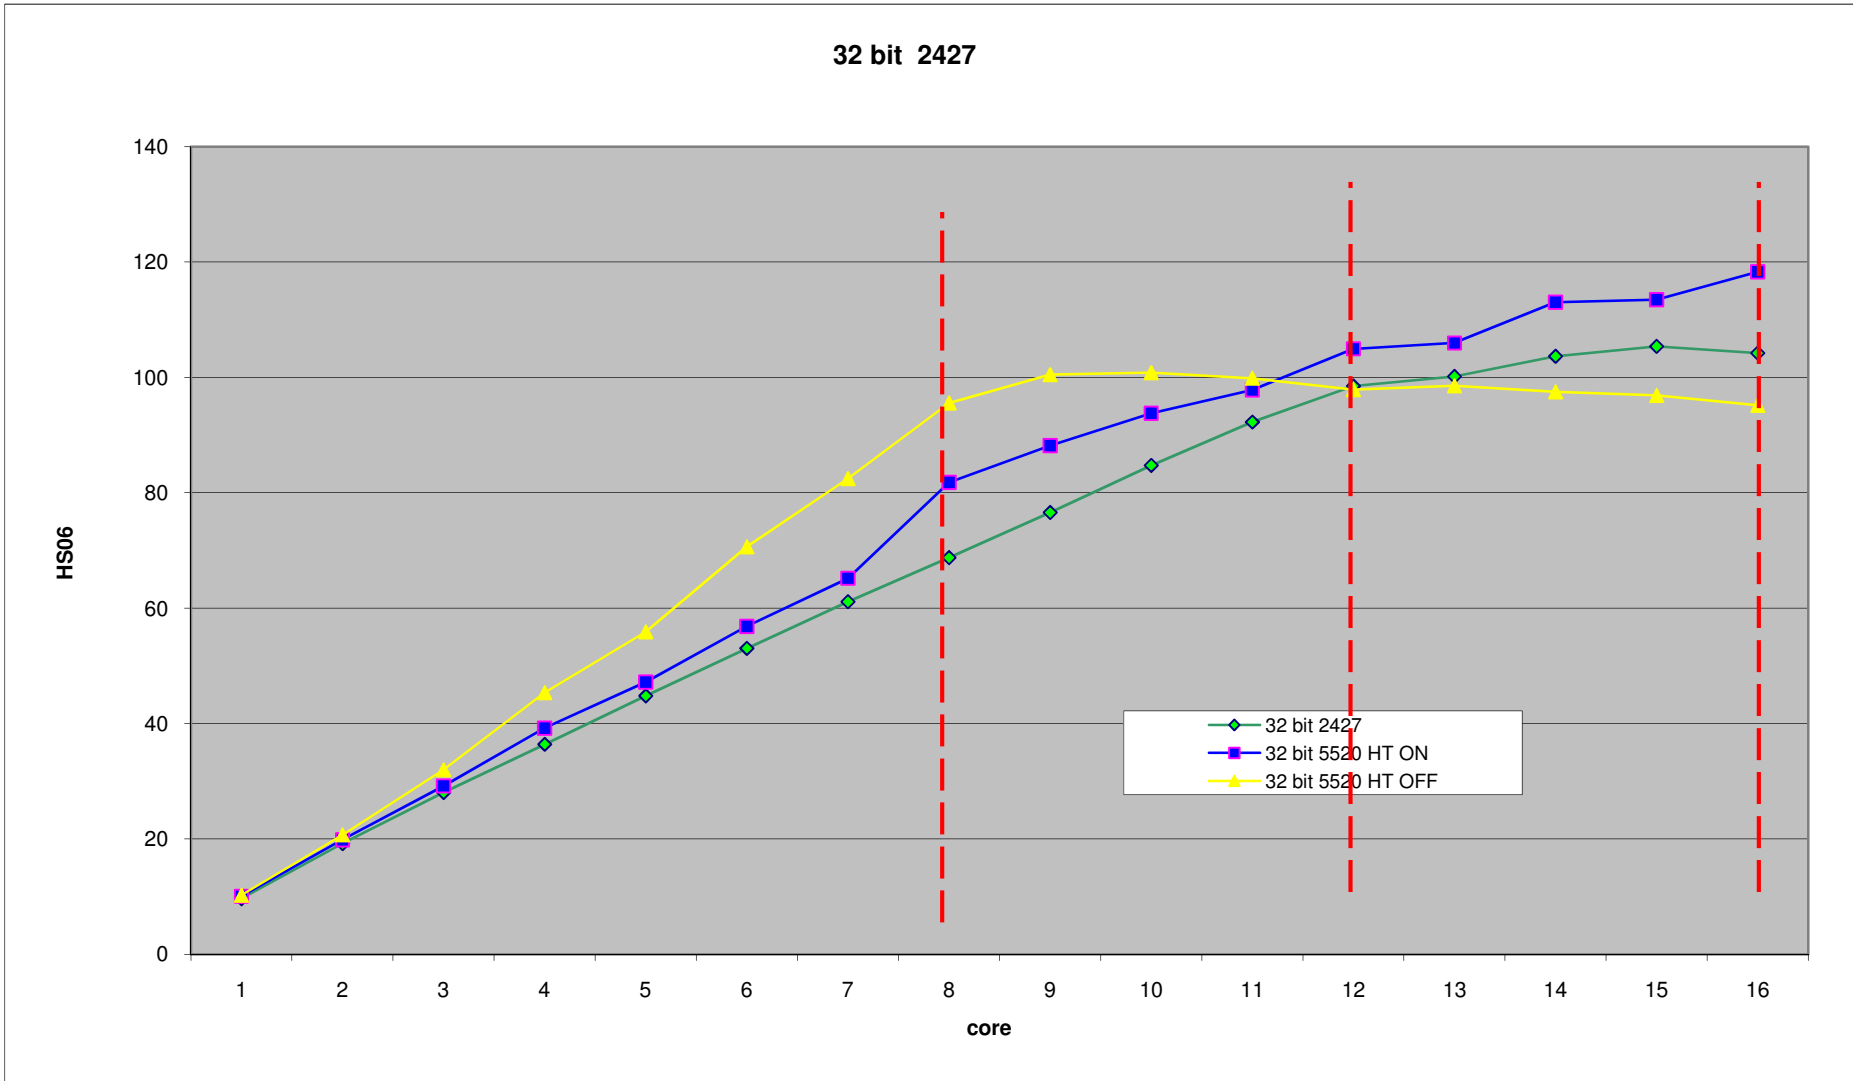





- Opteron 2427 – 2200 MHz – 32GB
- 2P x 6cores

- 5520 with HT OFF increased performance even after 8 cores
- 2247 doesn't show any drop with 12 cores fully loaded
- What happens if we start overloading with more processes than cores?

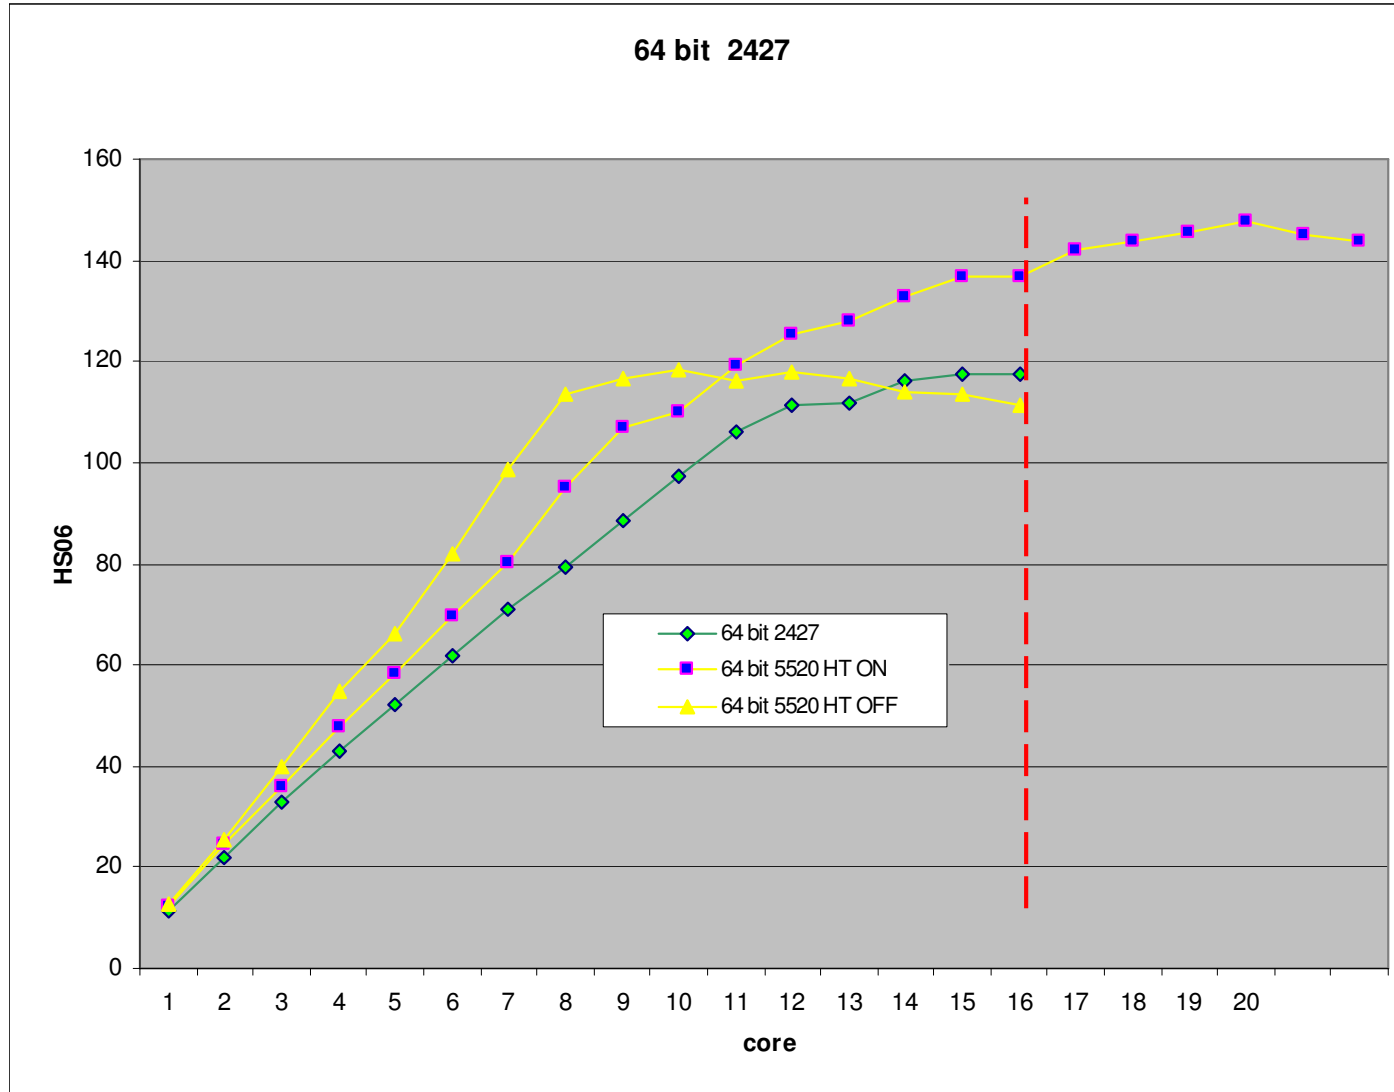
# 5520(2266) vs 2427(2200)



# 5520(2266) vs 2427(2200)



# 5520(2266) vs 2427(2200)

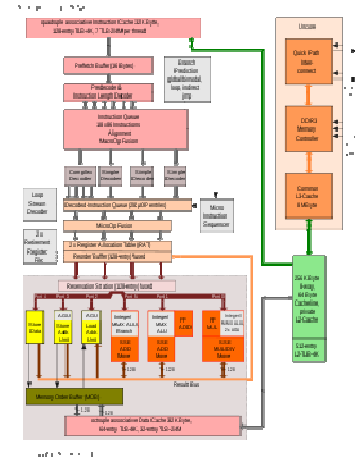


- Compare with Atlas and CMS code
  - GEN, SIM, DIGI and RECO
- Effect of HT, Turbo Mode and Overbooking on Power Consumption

# Questions?



- 45 nm
- Cache L1 32+32 KB
- Cache L2 256KB/core
- Cache L3 8MB shared
- 80W: E5502 1.86 GHz → E5540 2.53 GHz
- 95W: X5550 2.66 GHz → X5570 2.93 GHz
- Dual Thread (from 5520 upwards)
- Turbo Mode
- Quad core (excl. 5502 and 5508)





- 45 nm
- Cache L1 128 KB
- Cache L2 512KB/core
- Cache L3 6MB shared
- Power Consumption:
  - EE: max 40W (1.8 GHz)
  - HE: max 55W (2.0-2.1 GHz)
  - Standard max 75W (2.2-2.6 GHz)
  - SE max 105W (2.8GHz)

