CGG – TERATEC HPC HACKATON -ARM

Introduction to seismic imaging and stencil's algorithm

cgg.com

CGG

SEISMIC IMAGING

Why seismic imaging is needed?

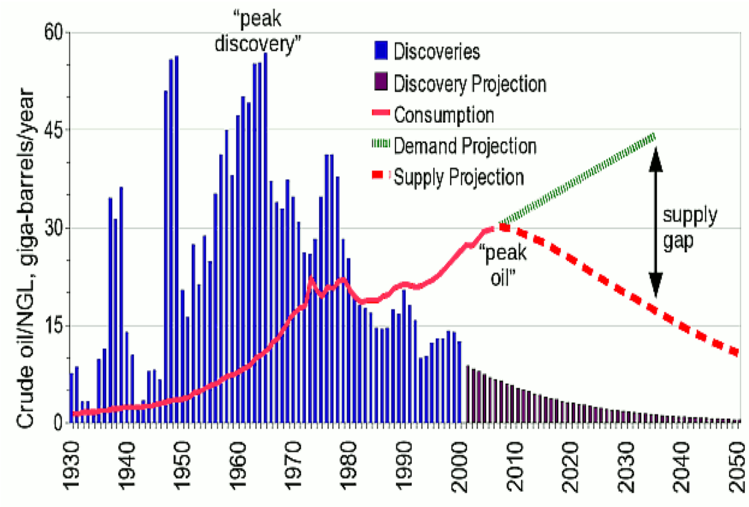

$ / Barrel	Investment Strategy
> $80	Drill, drill, drill
< $50	Optimize, Optimize, Optimize

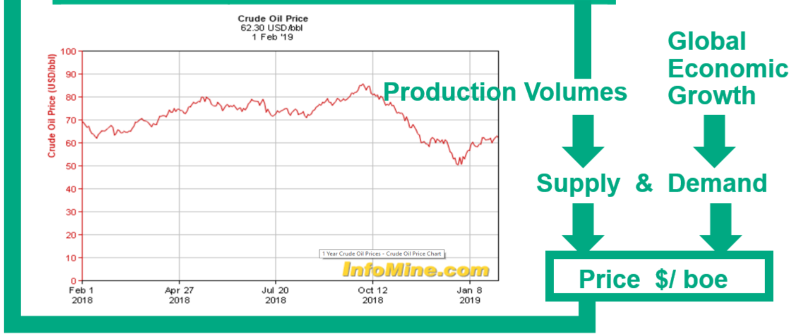

Necessity is the mother of inventions

HPC is a fundamental tool to help to improve discovery, minimize cost and risks.
Drilling price = 10x Seismic exploration and imaging price

Land and Marine acquisition: seismic surveys

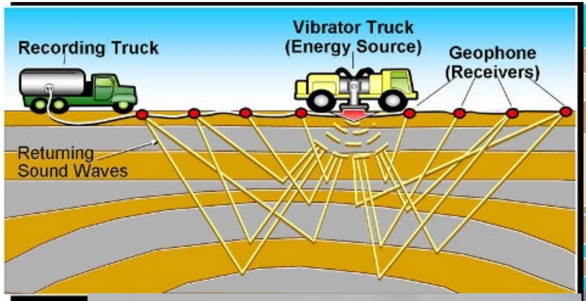

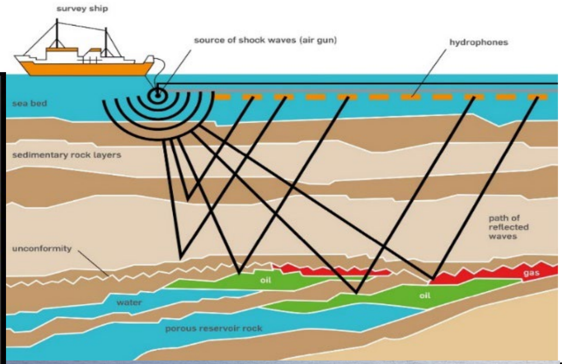

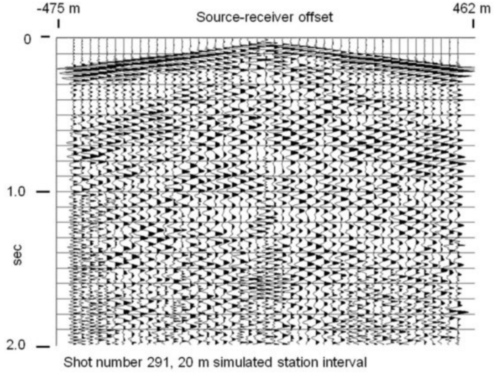

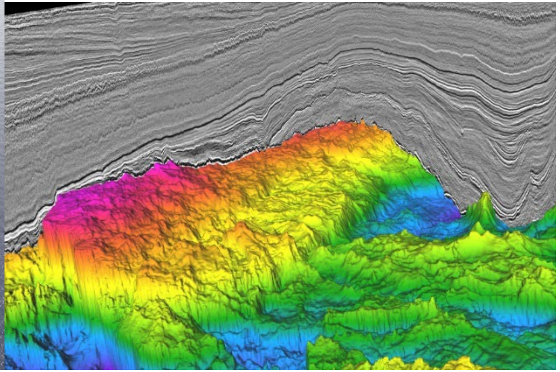

Offshore acquisition

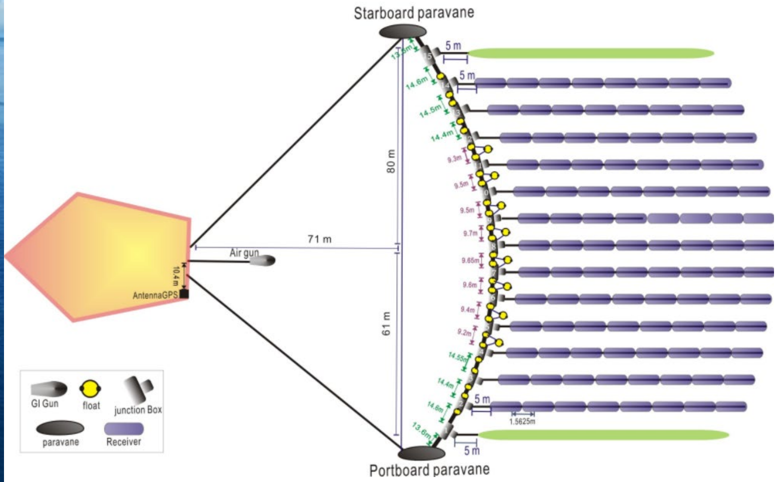

Because of the length of the cables (approx. 10km), the boat can't stop easily!

Onshore acquisition

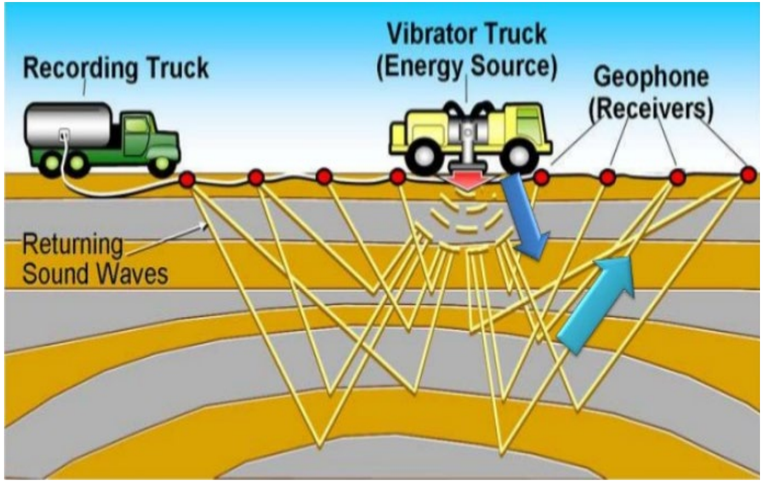

What is seismic IMAGING ?

1

Data acquisition: seismic surveys

2

Initial velocity/density model

Inversion algorithms FWI, RTM…

From the initial to the « best » velocity/density model

3

« Best » velocity/density model

Initial Velocity Model Building

Geologist → Where should I start looking ?

+

Experts → Restrict area and depth

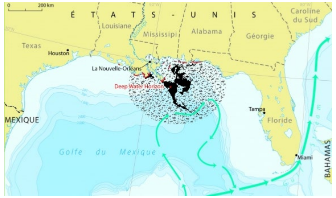

Deepwater horizon (BP) :
Water depth: 1259 meters
Vertical drilling depth: 10685 meters

+

Geophysicist → Physical parameters VP/VS/RHO → **Initial velocity/density model**

Why do I need an accurate velocity/density model ?

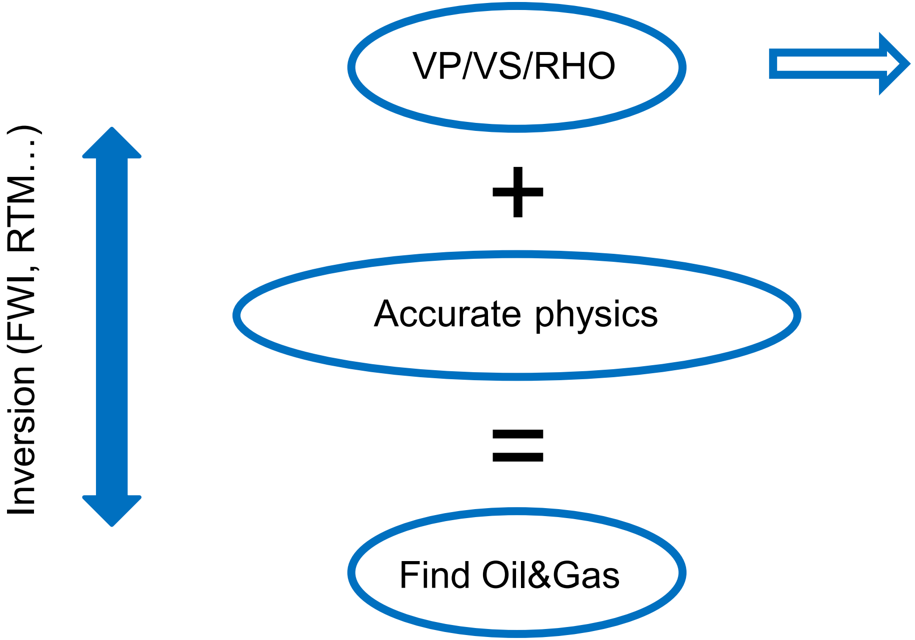

Type of formation	P wave velocity (m/s)	S wave velocity (m/s)	Density (g/cm^3)	Density of constituent crystal (g/cm^3)
Scree, vegetal soil	300-700	100-300	1.7-2.4	-
Dry sands	400-1200	100-500	1.5-1.7	2.65 quartz
Wet sands	1500-2000	400-600	1.9-2.1	2.65 quartz
Saturated shales and clays	1100-2500	200-800	2.0-2.4	-
Marls	2000-3000	750-1500	2.1-2.6	-
Saturated shale and sand sections	1500-2200	500-750	2.1-2.4	-
Porous and saturated sandstones	2000-3500	800-1800	2.1-2.4	2.65 quartz
Limestones	3500-6000	2000-3300	2.4-2.7	2.71 calcite
Chalk	2300-2600	1100-1300	1.8-3.1	2.71 calcite
Salt	4500-5500	2500-3100	2.1-2.3	2.1 halite
Anhydrite	4000-5500	2200-3100	2.9-3.0	-
Dolomite	3500-6500	1900-3600	2.5-2.9	(Ca, Mg) CO_3 2.8-2.9
Granite	4500-6000	2500-3300	2.5-2.7	-
Basalt	5000-6000	2800-3400	2.7-3.1	-
Gneiss	4400-5200	2700-3200	2.5-2.7	-
Coal	2200-2700	1000-1400	1.3-1.8	-
Water	1450-1500	-	1.0	-
Ice	3400-3800	1700-1900	0.9	-
Oil	1200-1250	-	0.6-0.9	-

Tell me your speed and I will tell you who you are…

Real vs Initial velocity/density model

Initial velocity/density model

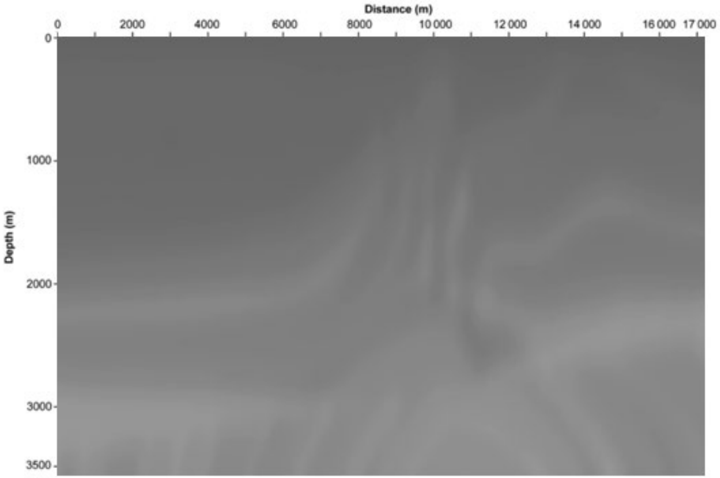

How do we get there ?

Real velocity/density model

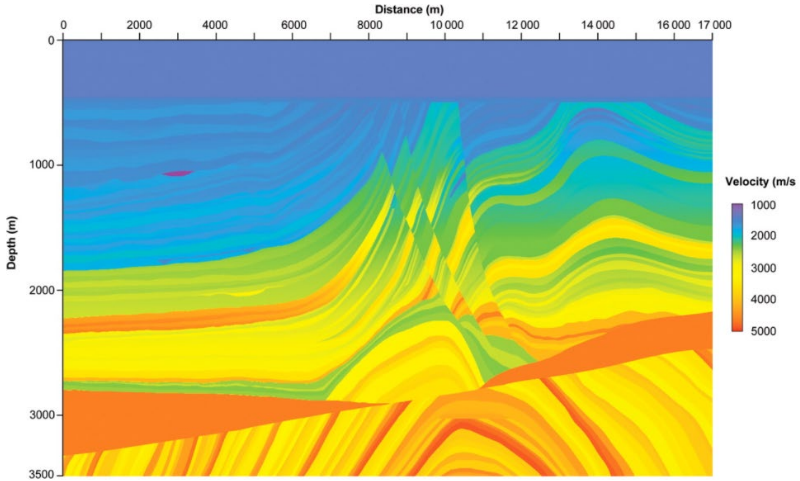

Inversion algorithms FWI, RTM…

Start with initial velocity/density model

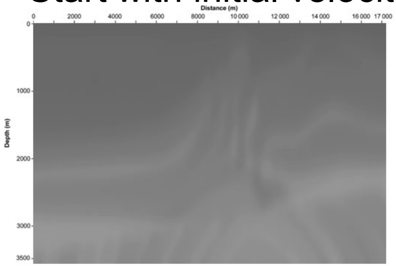

Propagate waves and record data « numerically »

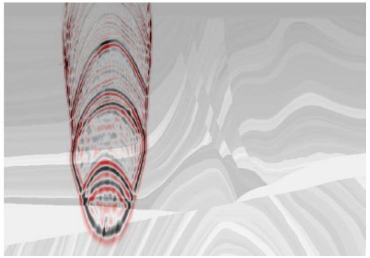

Compute differences between data acquisition (surveys) and numerical data
+ Compute gradient + update velocity/density model

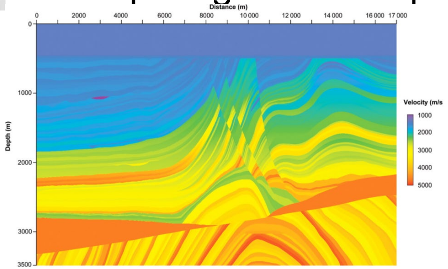

Exponential increase in compute required!

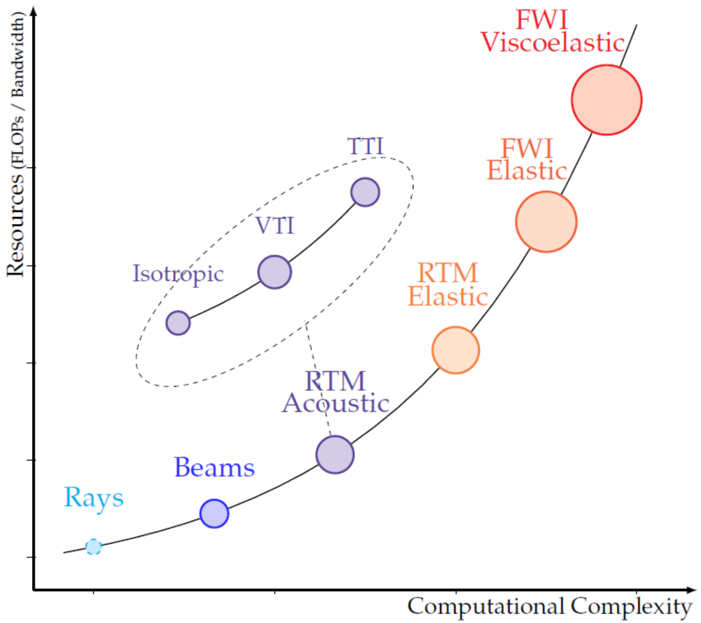

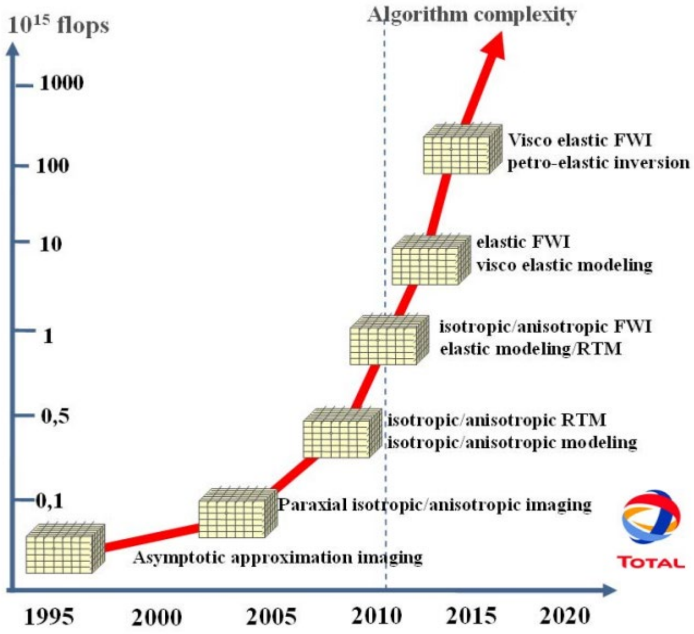

Algorithmic complexity and corresponding computing power

STENCIL

Stencil operators

Stencil Order	Extent	Memory Accesses/Elem.	Flops/Elem.
2	3×3×3	8	8
4	5×5×5	14	15
6	7×7×7	20	22
8	9×9×9	26	29
10	11×11×11	32	36
12	13×13×13	38	43

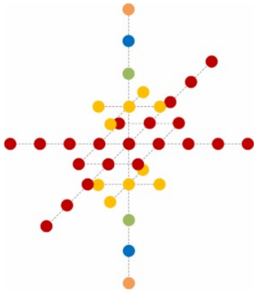

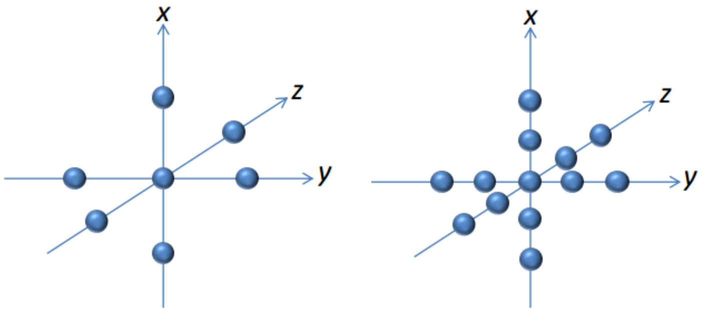

(a) 2nd order stencil (b) 4th order stencil

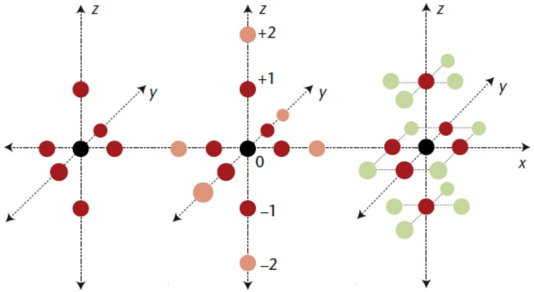

Stencils are workhorse in many HPC kernels, dominant in seismic industry

Pseudocode and modeling of a kernel

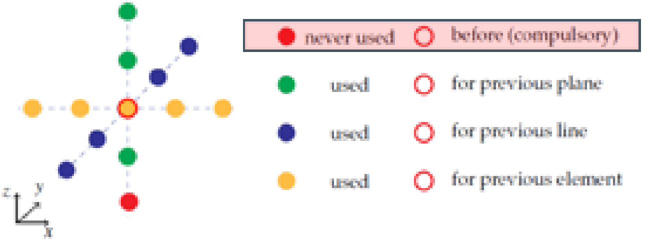

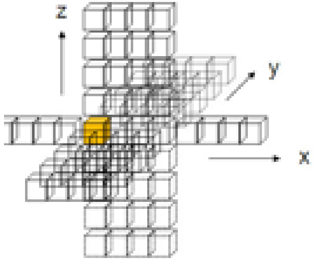


```
timesteps=1..1000
nz=1..800
ny=1..1200
nx=1..1600
order=2

float S[nx,ny,nz]   // Source       array
float D[nx,ny,nz]   // Destination  array
float C[nx,ny,nz]   // coefficients array
for t in [1..timesteps]
        for z in  [order+1..nz-order]
                for y in [order+1..ny-order]
                        for x in [order+1..nx-order]
                          D[x, y, z] = S[ x    , y    , z+1 ]*coeff1 + S[ x    , y    , z+2 ]*coeff2
                                     + S[ x    , y    , z-1 ]*coeff1 + S[ x    , y    , z-2 ]*coeff2
                                     + S[ x    , y+1  , z   ]*coeff1 + S[ x    , y+2  , z   ]*coeff2
                                     + S[ x    , y-1  , z   ]*coeff1 + S[ x    , y-2  , z   ]*coeff2
                                     + S[ x+1  , y    , z   ]*coeff1 + S[ x+2  , y    , z   ]*coeff2
                                     + S[ x-1  , y    , z   ]*coeff1 + S[ x-2  , y    , z   ]*coeff2
                                     + S[ x    , y    , z   ]*coeff0 + C[ x    , y    , z   ]
                        endfor x //  13 MUL SP + 13 ADD SP ; can use 13 FMA
                                 //  if well implemented only RED is compulsory then comes from DRAM
                                 //  Green and Blue come from recent accesses; L1+L2 and LLC if large enough
                 endfor y
        endfor z
        swap(S&,D&)
endfor t

// per point : 26 flops ; 4 bytes C ; 4 bytes S ; 4 bytes D = 16 bytes on DRAM bus : 11 streams
// kernel is memory bound ; flops/bytes ratio = 26/12 = 2.2
```

Roofline model and academic experiments

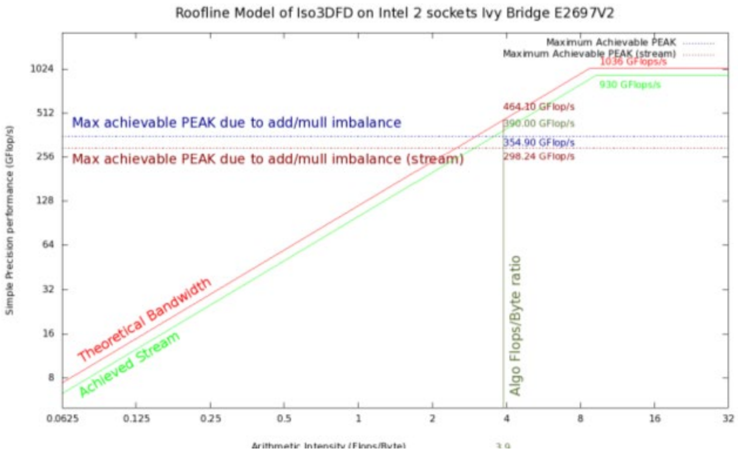

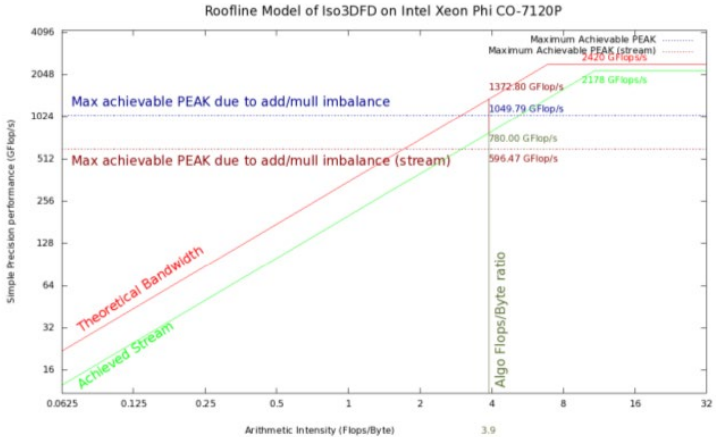

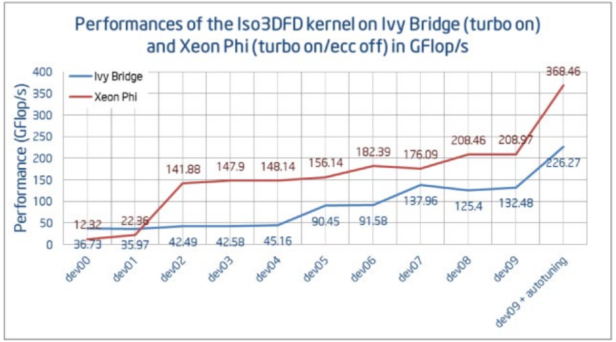


```
IVY : peak 1036 GFlops : effective 226 Gflops
Ratio peak/effective = 5 =   20% efficiency
 KNC : peak 2420 Gflops : effective 368 Gflops
Ratio peak/effective = 6.6 = 15% efficiency
Stencils are memory bound
```

http://www.techenablement.com/characterization-optimization-methodology-applied-stencil-computations/

Trend in Gflops and Bandwidth of latest Xeon

Memory bandwidth vs Gflops imbalance continue to grow

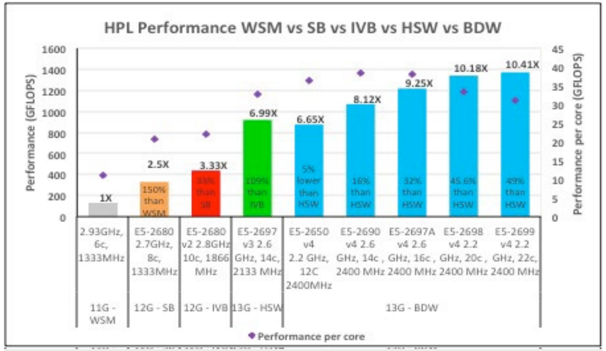

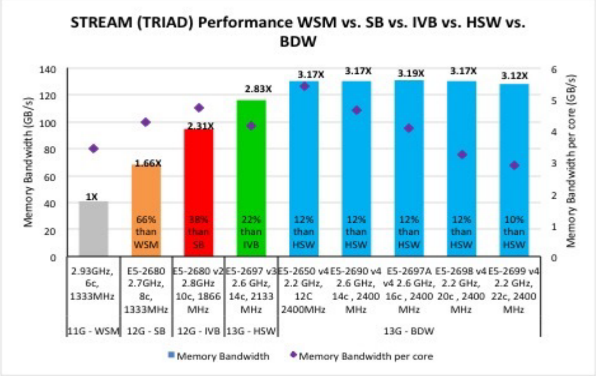

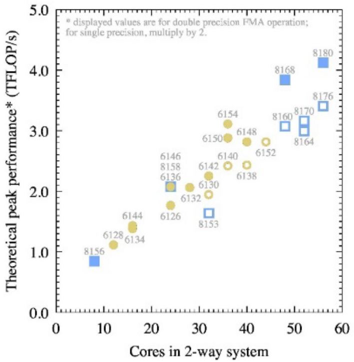

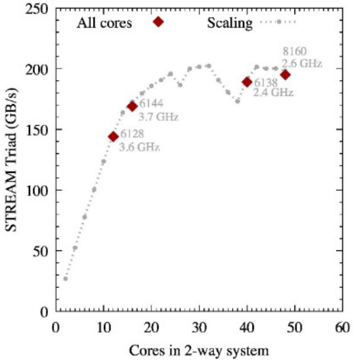

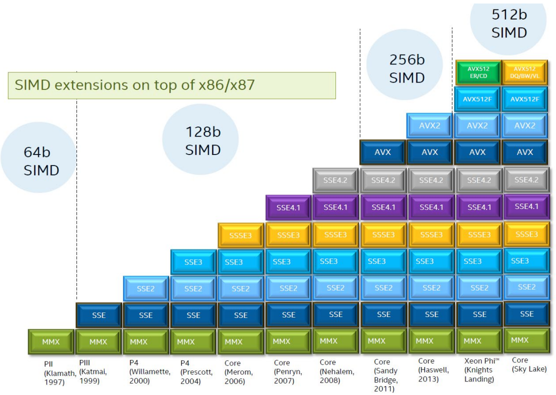

Names	Memory Clock	I/O Bus Clock	Transfer Rate	Theoretical Bandwidth
DDR-200, PC-1600	100 MHz	100 MHz	0.2 GT/s	1.6 GB/s
DDR2-800, PC2-6400	200 MHz	400 MHz	0.8 GT/s	6.4 GB/s
DDR3-1600, PC3-12800	200 MHz	800 MHz	1.6 GT/s	12.8 GB/s
DDR4-3200, PC4-25600	400 MHz	1600 MHz	3.2 GT/s	25.6 GB/s

https://colfaxresearch.com/xeon-2017/

Skylake socket

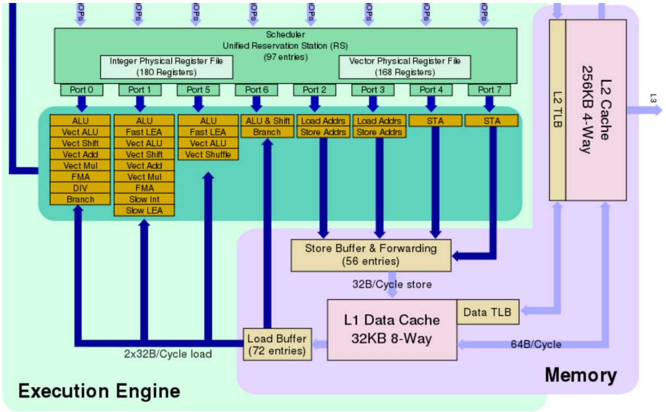

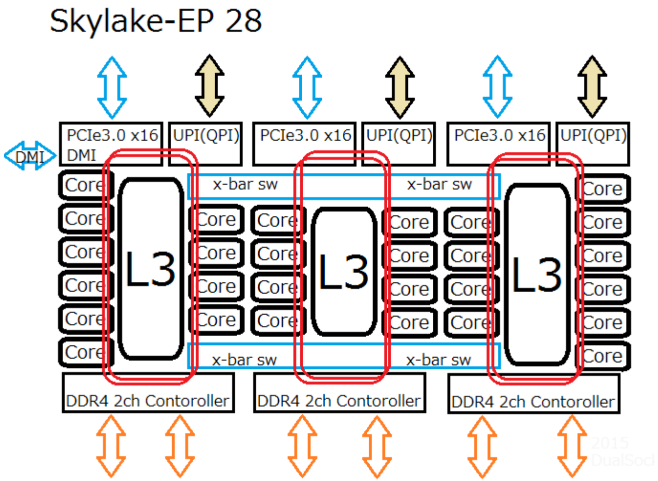

AVX 512 = 2 FMA/cycle = 64 SP/cycle
28 cores = 2.6 Tflops SP at 2.0Ghz
L1 and L2 can deliver 128 GB/s at 2.0 Ghz

Bandwidth DRAM =
64bits*2400MTS*6channels/8bits

Effective 90% ~ 100 GB/s per socket

Need 26 flops per byte to cover the memory bandwidth

Options for multicores architectures

DDR4/3 28nm	HBM2 14nm	HBM3 7nm	DDR5 7nm
• 3200Mbps • x16 – x72-bits • 1-4 Ranks • DFI 4.0	• 2000Mbps • 1024-bit • 2.5D design architecture	• Expected 4000Mbps • Complex design architectures	• Expected 4800 – 6400Mbps

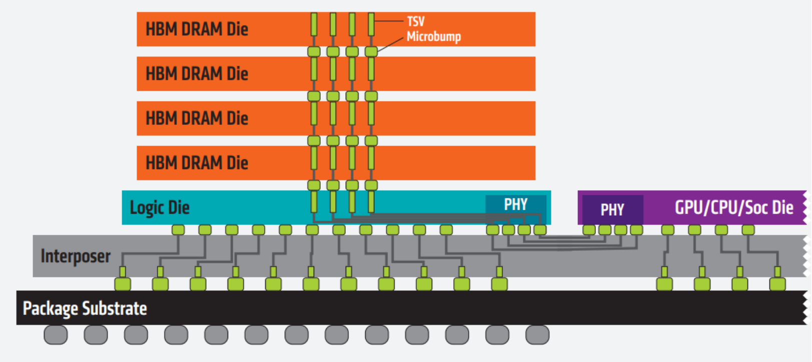

- DDR5 and 8 channels per SOC will deliver 250GB/s/SOC
- HBM2 can deliver 300GB/S per stack
- HBM3 can deliver 500GB/s per stack
- High capacity with SCM
- O&G need
  - High bandwidth to cover flops
  - High capacity for IO in memory

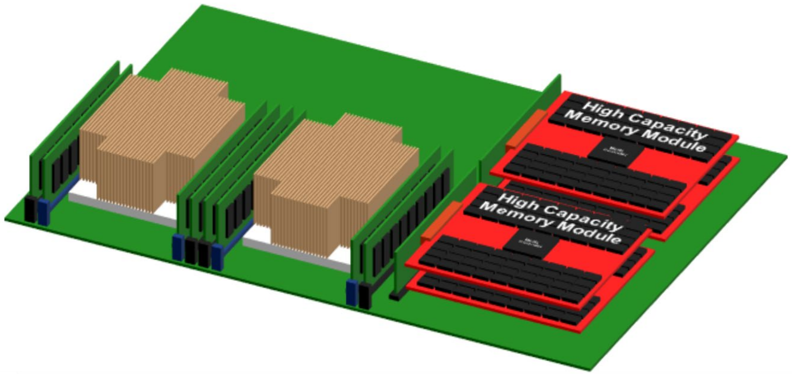

Thank you, and good luck for ARM porting!