

High Performance SAR Image Formation On Commodity Multicore Architectures

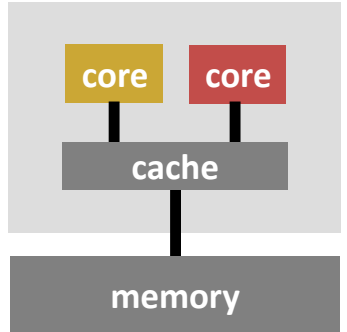
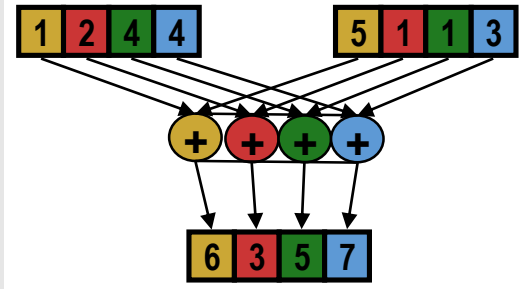
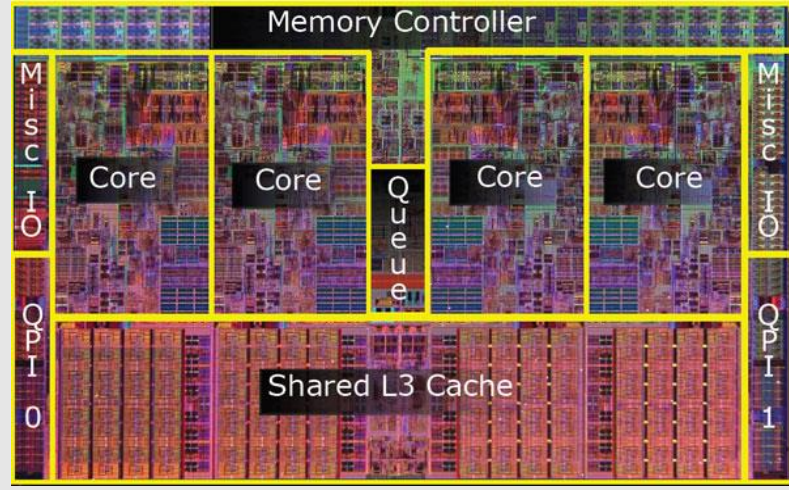
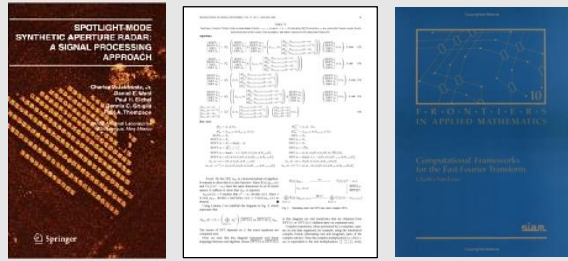
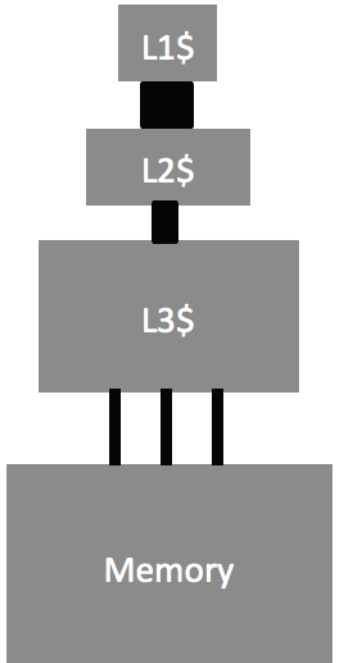
Daniel S. McFarlin

Franz Franchetti

Markus Püschel

José M. F. Moura

Challenge: High Performance Mapping of Algorithms to Highly Parallel Hardware

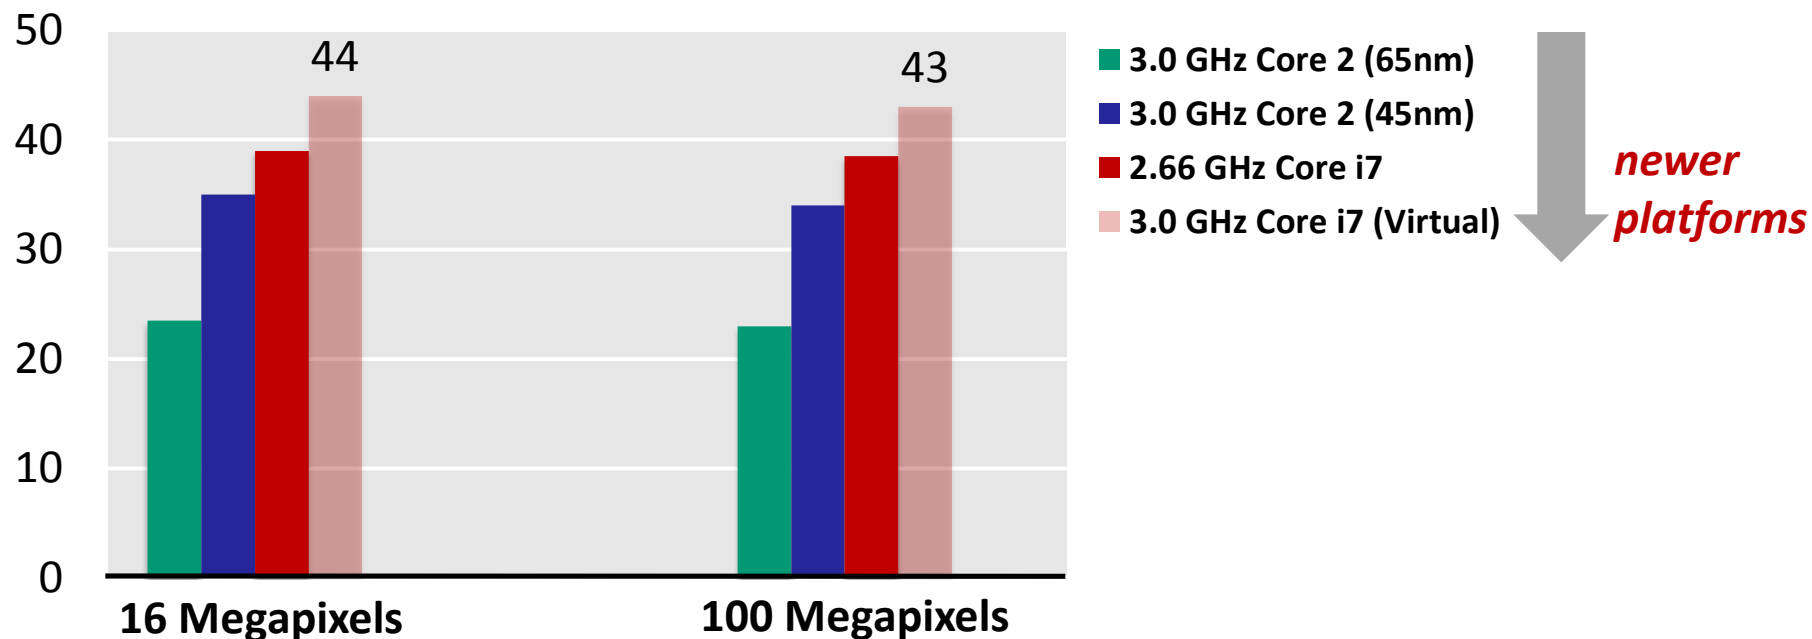


Difficult for General Purpose C Toolchain

Spiral's Automatically Generated PFA SAR Image Formation Code

SAR Image Formation on Intel platforms

performance [Gflop/s]



- Algorithm by J. Rudin (best paper award, HPEC 2007): 30 Gflop/s on Cell
- Each implementation: vectorized, threaded, cache tuned, ~13 MB of code
- *Code was not written by a human*

Transform
user specified

DFT₈



Fast algorithm
in SPL

$$(DFT_2 \otimes I_4) T_4^8 (I_2 \otimes ((DFT_2 \otimes I_2) \cdot T_2^4 (I_2 \otimes DFT_2) L_2^4)) L_2^8$$

many choices



Σ-SPL:

$$\sum (S_j DFT_2 G_j) \sum \left(\sum (S_{k,l} \text{diag}(t_{k,l}) DFT_2 G_l) \sum (S_m \text{diag}(t_m) DFT_2 G_{k,m}) \right)$$



C Code:

```
void sub(double *y, double *x) {  
    double f0, f1, f2, f3, f4, f7, f8, f10, f11;  
    f0 = x[0] - x[3];  
    f1 = x[0] + x[3];  
    f2 = x[1] - x[2];  
    f3 = x[1] + x[2];  
    f4 = f1 - f3;  
    y[0] = f1 + f3;  
    y[2] = 0.7071067811865476 * f4;  
    f7 = 0.9238795325112867 * f0;  
    f8 = 0.3826834323650898 * f2;  
    y[1] = f7 + f8;  
    f10 = 0.3826834323650898 * f0;  
    f11 = (-0.9238795325112867) * f2;  
    y[3] = f10 + f11;  
}
```

**Optimization at all
abstraction levels**



**parallelization
vectorization**



**loop
optimizations**

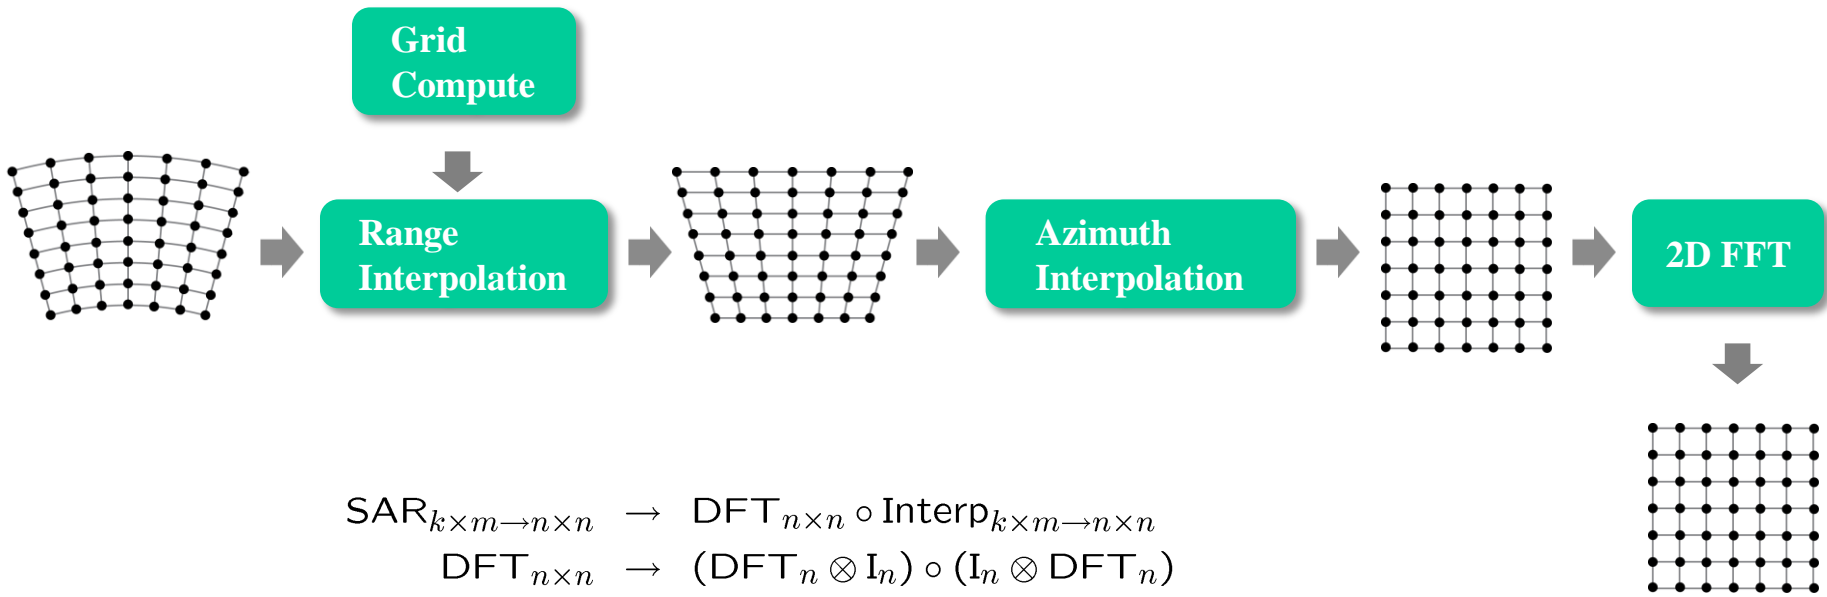


**constant folding
scheduling**

.....

Iteration of this process to search for the fastest

Spiral Formula Representation of SAR



$$\begin{aligned}
 \text{SAR}_{k \times m \rightarrow n \times n} &\rightarrow \text{DFT}_{n \times n} \circ \text{Interp}_{k \times m \rightarrow n \times n} \\
 \text{DFT}_{n \times n} &\rightarrow (\text{DFT}_n \otimes \mathbf{I}_n) \circ (\mathbf{I}_n \otimes \text{DFT}_n) \\
 \text{Interp}_{k \times m \rightarrow n \times n} &\rightarrow (\text{Interp}_{k \rightarrow n} \otimes_i \mathbf{I}_n) \circ (\mathbf{I}_k \otimes_i \text{Interp}_{m \rightarrow n}) \\
 \text{Interp}_{r \rightarrow s} &\rightarrow \left(\bigoplus_{i=0}^{n-2} \text{InterpSeg}_k \right) \oplus \text{InterpSegPruned}_{k,l} \\
 \text{InterpSeg}_k &\rightarrow G_f^{u \cdot n \rightarrow k} \circ \text{iPrunedDFT}_{n \rightarrow u \cdot n} \circ \left(\frac{1}{n} \right) \circ \text{DFT}_n
 \end{aligned}$$

Parallelization through Rewriting

Threading:

$$\begin{aligned}
 \underbrace{\text{DFT}_{mn}}_{\text{smp}(p,\mu)} &\rightarrow \underbrace{((\text{DFT}_m \otimes I_n) \tau_n^{mn} (I_m \otimes \text{DFT}_n) L_m^{mn})}_{\text{smp}(p,\mu)} \\
 &\dots \\
 &\rightarrow \underbrace{(\text{DFT}_m \otimes I_n)}_{\text{smp}(p,\mu)} \underbrace{\tau_n^{mn}}_{\text{smp}(p,\mu)} \underbrace{(I_m \otimes \text{DFT}_n)}_{\text{smp}(p,\mu)} \underbrace{L_m^{mn}}_{\text{smp}(p,\mu)} \\
 &\dots \\
 &\rightarrow ((L_m^{mp} \otimes I_{n/p\mu}) \otimes_{\mu} I_{\mu}) (I_p \otimes_{\parallel} (\text{DFT}_m \otimes I_{n/p})) ((L_p^{mp} \otimes I_{n/p\mu}) \otimes_{\mu} I_{\mu}) \\
 &\quad \left(\bigoplus_{i=0}^{p-1} \tau_n^{mn,i} \right) (I_p \otimes_{\parallel} (I_{m/p} \otimes \text{DFT}_n)) (I_p \otimes_{\parallel} L_{m/p}^{mn/p}) ((L_p^{pn} \otimes I_{m/p\mu}) \otimes_{\mu} I_{\mu})
 \end{aligned}$$

Vectorization:

$$\begin{aligned}
 \underbrace{(\text{DFT}_{mn})}_{\text{vec}(\nu)} &\rightarrow \underbrace{((\text{DFT}_m \otimes I_n) \tau_n^{mn} (I_m \otimes \text{DFT}_n) L_m^{mn})}_{\text{vec}(\nu)} \\
 &\dots \\
 &\rightarrow \underbrace{(\text{DFT}_m \otimes I_n)^{\nu}}_{\text{vec}(\nu)} \underbrace{(\tau_n^{mn})^{\nu}}_{\text{vec}(\nu)} \underbrace{(I_m \otimes \text{DFT}_n) L_m^{mn}}_{\text{vec}(\nu)} \\
 &\dots \\
 &\rightarrow (I_{mn/\nu} \otimes \underbrace{L_{\nu}^{2\nu}}_{\text{sse}}) (\overline{\text{DFT}_m \otimes I_{n/\nu}} \otimes I_{\nu}) \underbrace{(\tau_n^{mn})^{\nu}}_{\text{sse}} \\
 &\quad (I_{m/\nu} \otimes (\overline{L_{\nu}^n} \otimes I_{\nu})) (I_{n/\nu} \otimes (\overline{L_{\nu}^{2\nu}} \otimes I_{\nu})) (I_2 \otimes \underbrace{L_{\nu}^{\nu^2}}_{\text{sse}}) (\overline{L_{\nu}^{2\nu}} \otimes I_{\nu}) (\overline{\text{DFT}_n} \otimes I_{\nu}) \\
 &\quad ((L_m^{mn} \otimes I_2) \otimes I_{\nu}) (I_{mn/\nu} \otimes \underbrace{L_{\nu}^{2\nu}}_{\text{sse}})
 \end{aligned}$$

GPUs:

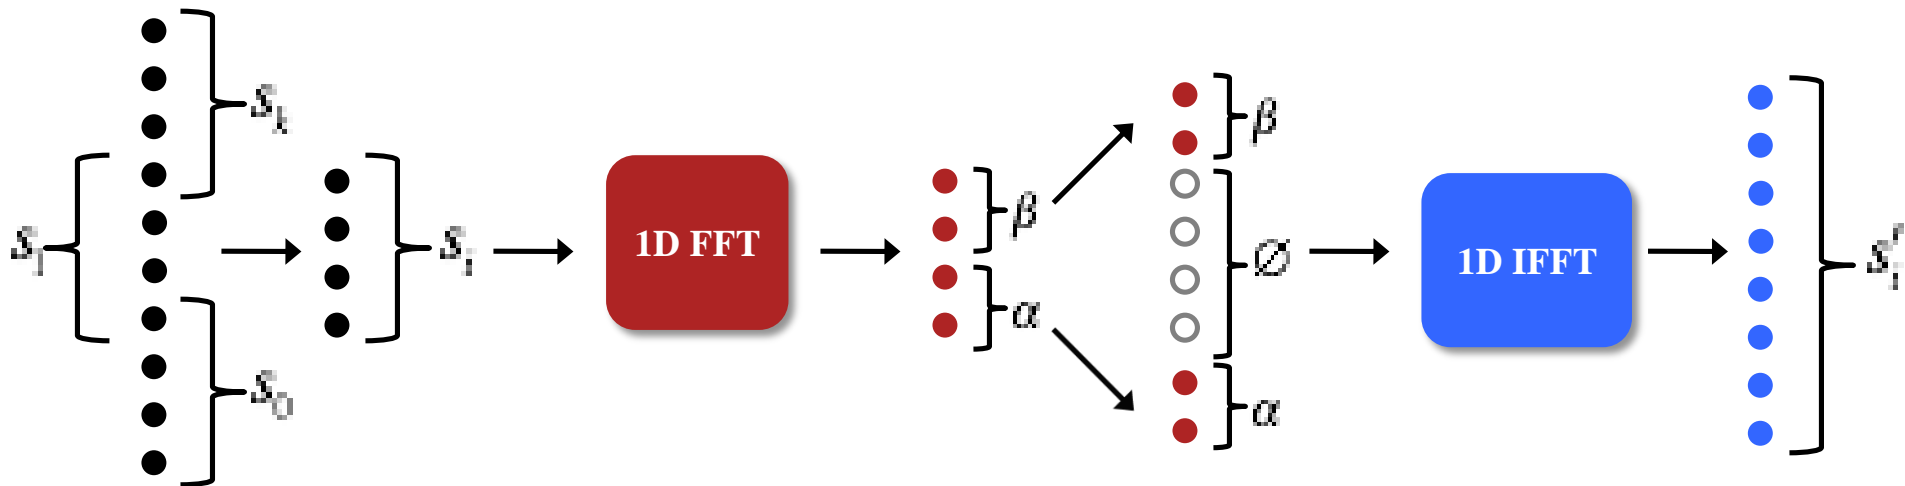
$$\begin{aligned}
 \underbrace{(\text{DFT}_{r^k})}_{\text{gpu}(t,c)} &\rightarrow \underbrace{\left(\prod_{i=0}^{k-1} L_r^{r^k} (I_{r^{k-1}} \otimes \text{DFT}_r) (L_{r^{k-i-1}} (I_{r^i} \otimes \tau_{r^{k-i-1}}) L_{r^{i+1}}^{r^k}) \right)}_{\text{gpu}(t,c)} R_r^{r^k} \\
 &\dots \\
 &\rightarrow \left(\prod_{i=0}^{k-1} (L_r^{r^n/2} \otimes I_2) (I_{r^{n-1}/2} \otimes \times \underbrace{(\text{DFT}_r \otimes I_2) L_r^{2r}}_{\text{shd}(t,c)} \tau_i) \right) \\
 &\quad (L_r^{r^n/2} \otimes I_2) (I_{r^{n-1}/2} \otimes \times \underbrace{L_r^{2r}}_{\text{shd}(t,c)}) (R_r^{r^{n-1}} \otimes I_r)
 \end{aligned}$$

Verilog for FPGAs:

$$\begin{aligned}
 \underbrace{(\text{DFT}_{r^k})}_{\text{stream}(r^s)} &\rightarrow \underbrace{\left[\prod_{i=0}^{k-1} L_r^{r^k} (I_{r^{k-1}} \otimes \text{DFT}_r) (L_{r^{k-i-1}} (I_{r^i} \otimes \tau_{r^{k-i-1}}) L_{r^{i+1}}^{r^k}) \right]}_{\text{stream}(r^s)} R_r^{r^k} \\
 &\dots \\
 &\rightarrow \left[\prod_{i=0}^{k-1} \underbrace{L_r^{r^k}}_{\text{stream}(r^s)} \underbrace{(I_{r^{k-1}} \otimes \text{DFT}_r)}_{\text{stream}(r^s)} \underbrace{(L_{r^{k-i-1}} (I_{r^i} \otimes \tau_{r^{k-i-1}}) L_{r^{i+1}}^{r^k})}_{\text{stream}(r^s)} \right] R_r^{r^k} \\
 &\dots \\
 &\rightarrow \left[\prod_{i=0}^{k-1} \underbrace{L_r^{r^k}}_{\text{stream}(r^s)} (I_{r^{k-s-1}} \otimes_s (I_{r^{s-1}} \otimes \text{DFT}_r)) \underbrace{\tau_i'}_{\text{stream}(r^s)} \right] R_r^{r^k} \\
 &\quad \text{stream}(r^s)
 \end{aligned}$$

- Rigorous, correct by construction
- Overcomes compiler limitations

Domain Specific FFT



- **Segmented Interpolation**

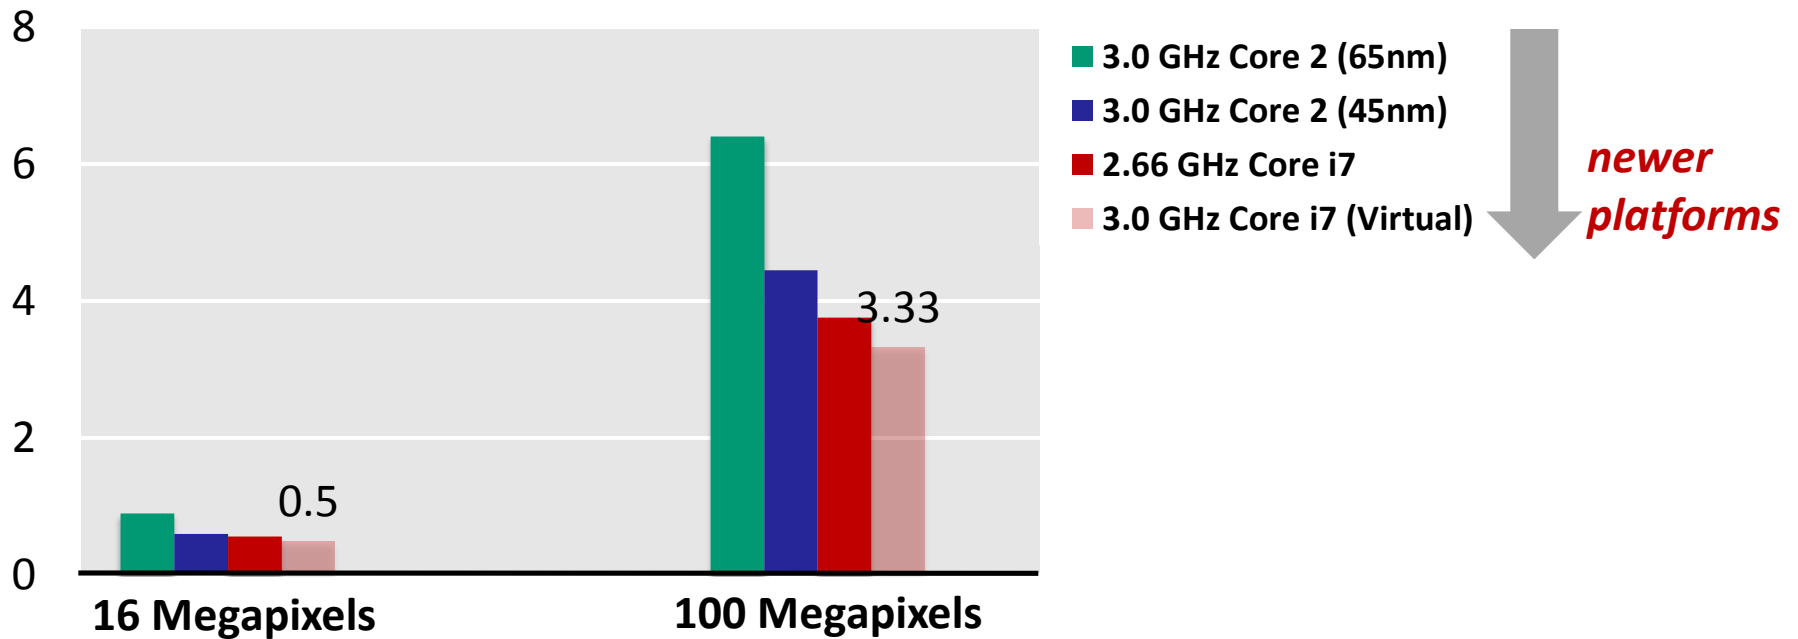
k segments of length *r*, with *u*-fold upsampling

Pruned FFT can reduce dominant Interpolation opcount by up to 15%

Performance Results

SAR Image Formation on Intel platforms

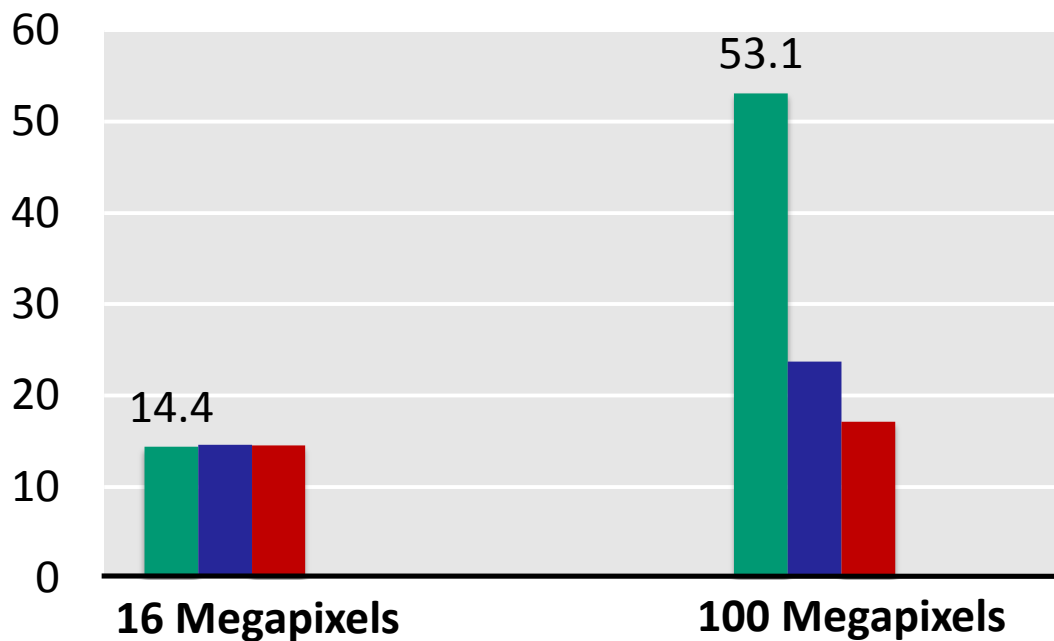
runtime [sec]



Performance Results

SAR Image Formation on Intel platforms

Percentage speedup of 2 MB pages over 4K pages



■ 3.0 GHz Core 2 (65nm)

■ 3.0 GHz Core 2 (45nm)

■ 2.66 GHz Core i7

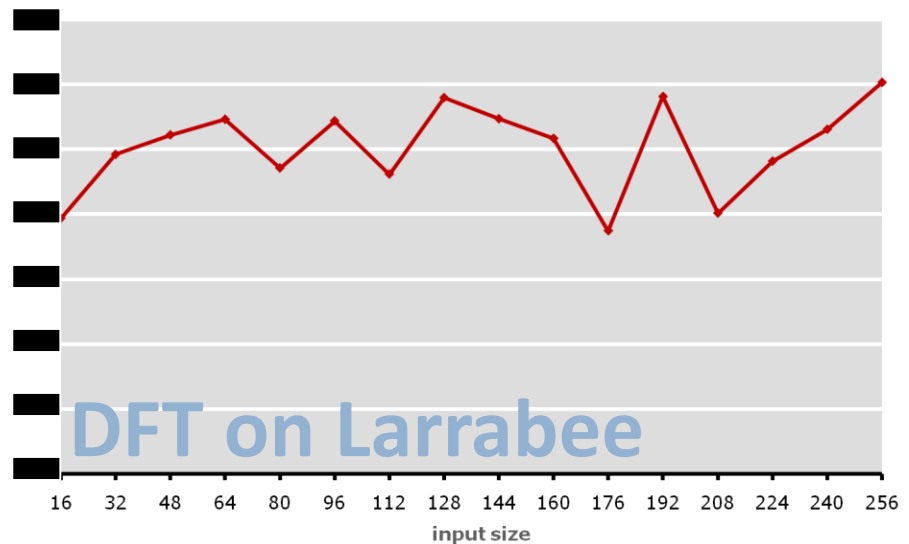


*newer
platforms*

Conclusions and Future Work

- Spiral generated SAR Image Formation performance comparable to hand-tuned code on the Cell
- SAR generation for non-released platforms
 - AVX
 - Larrabee

DFT (double-precision), single core Larrabee Native
performance [Gflop/s]



Not actual data (NDA)