

Computer Generation of Platform-Adapted Physical Layer Software

Yevgen Voronenko (SpiralGen, CMU)

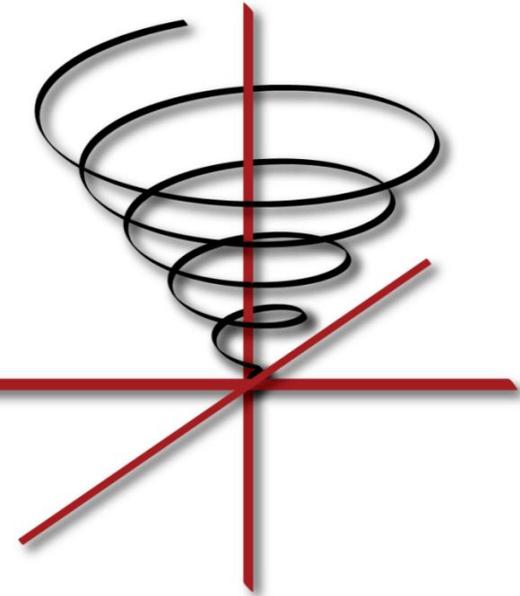
Christian Berger (CMU)

Volodymyr Arbatov (CMU)

Ronghui Peng (SpiralGen)

Franz Franchetti (CMU)

Markus Püschel (CMU)



SpiralGen at a Glance

■ Software & Service

For computationally intensive apps

Performance optimization service | Optimized components |
Performance evaluation

■ Value proposition

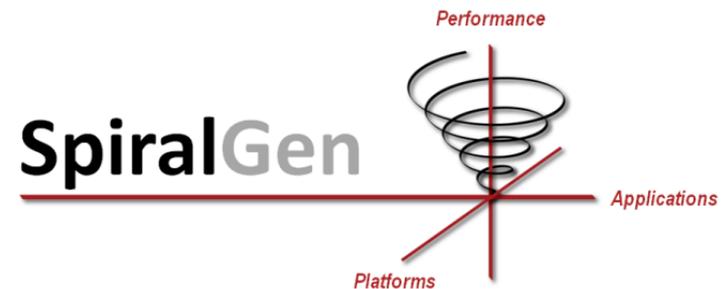
- *Faster time-to-market*
- *Quick customization*
- *Platform independence w/ no penalty*

■ Target applications:

- Core signal/image/video processing
- Medical Imaging

■ Target platforms:

- Multi-core DSPs / GPPs



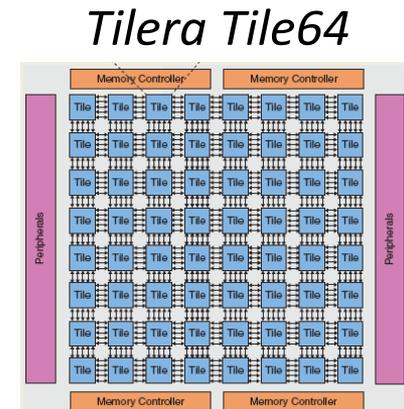
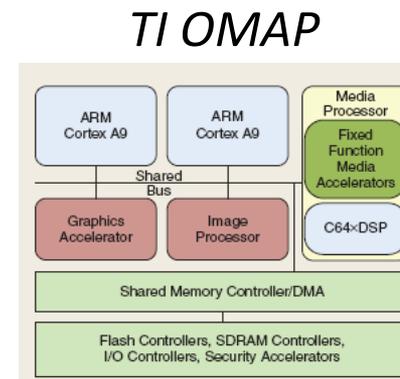
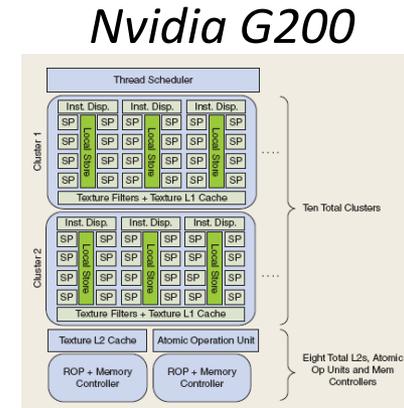
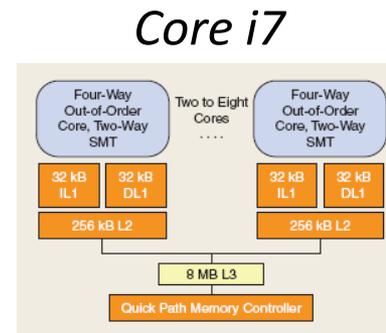
The Future is Parallel

■ General purpose/high performance platforms

- Intel Core i7 (4-8 cores, 128-B SIMD)
- Intel Larrabee (16, 512-B)
- ARM Cortex (1-4, -)
- IBM Cell (9, 128-B)
- Nvidia G200 (240)

■ DSP software platforms

- Tileria Tile64 (64, 3-way)
- Sandbridge SB3500 (4, 128-B)
- Freescale MSC8156 (6, ?)
- ~~Clearspeed~~
- ~~StreamProcessors~~

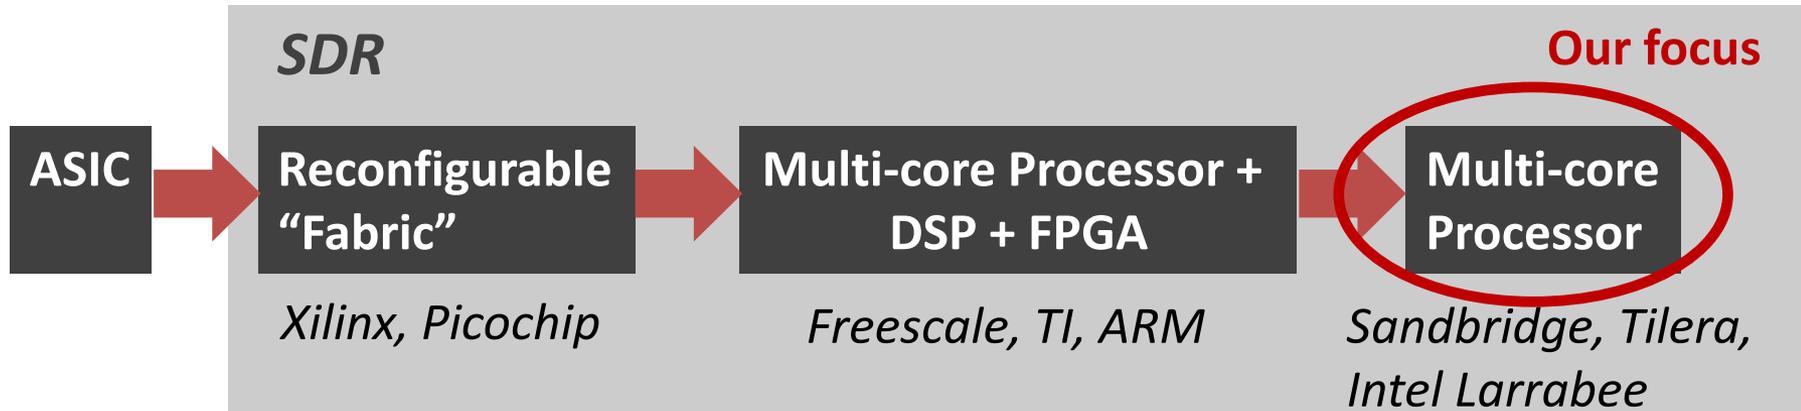


Source: IEEE SP Magazine, Vol. 26, November 2009

Can waveform porting be substantially accelerated?

Rethinking the Software Radio

- Trend for performance-driven apps?



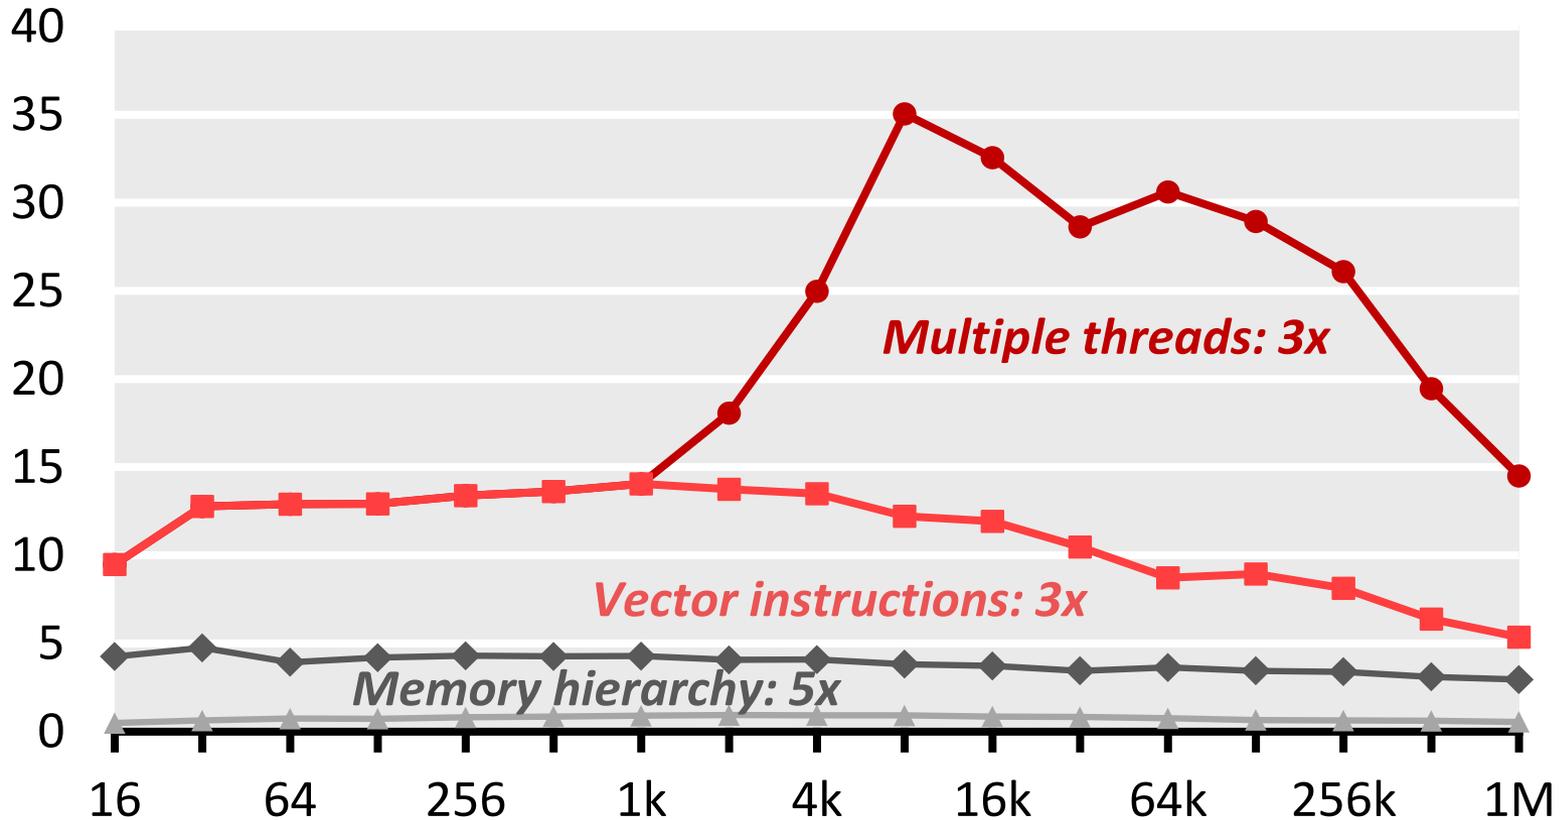
- **Key driver: reducing the chip count on board**
- **Advantages:**
 - Reduced software partitioning effort + easier porting
 - Lower power + longer battery life
- **Magic bullet: single-chip multi-layer SDR solution**
- **Example: DSP/GPP/FPGA board – 25W, Intel board – 30W**

Multi-core GPPs are real. But can compilers retarget the code?

The Answer is No

DFT (single precision) on Intel Core i7 (4 cores, 2.66 GHz)

Performance [Gflop/s]

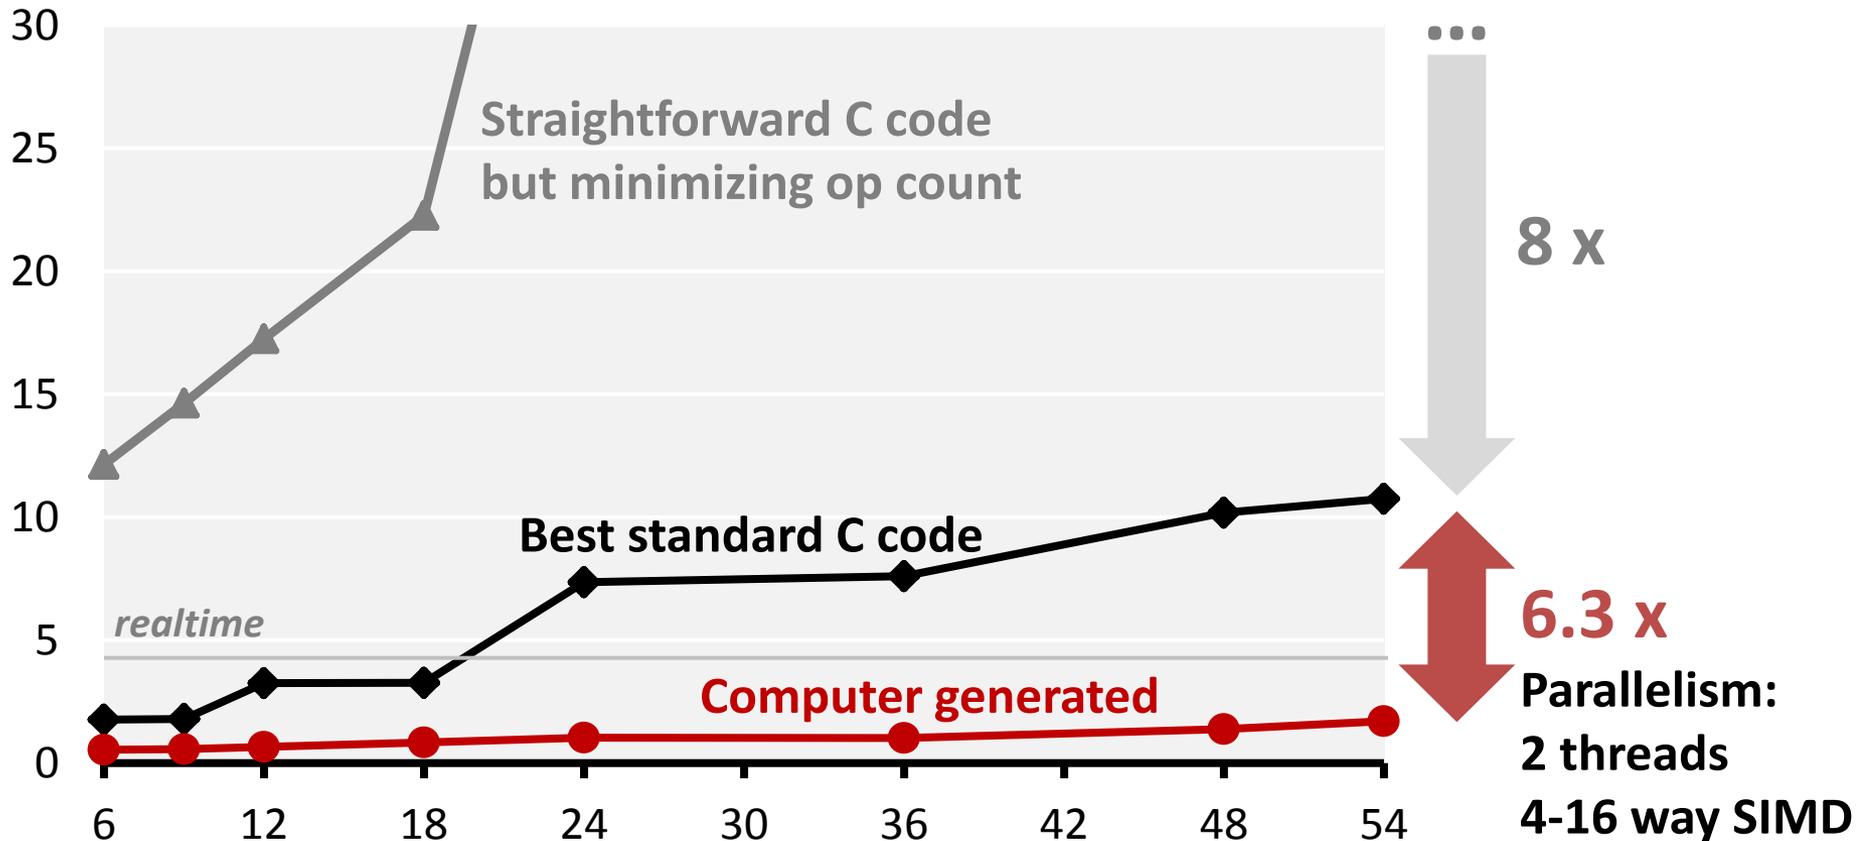


Only guru programmers can do the job

WiFi Transmitter Example

WiFi transmitter on Dualcore Intel Atom

Run time per OFDM symbol [micro seconds] vs. Data rate [Mbit/s]



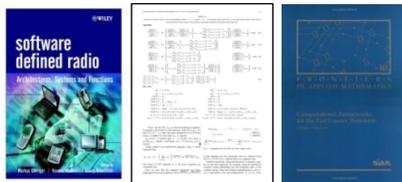
Compilers fail to optimize: 50x

SpiralGen's Core Technology:

Automation Tools

Application specification

Algorithm knowledge



Platform description

Instruction set architecture



```
_mm_set1_epi8(x) = ...  
_mm_xor_si128(x,y) = ...  
_mm_avg_epu8(x,y) = ...  
_mm_cmpeq_epi8(x,y) = ...  
_mm_unpacklo_epi8(x,y) = ...  
...
```



Spiral

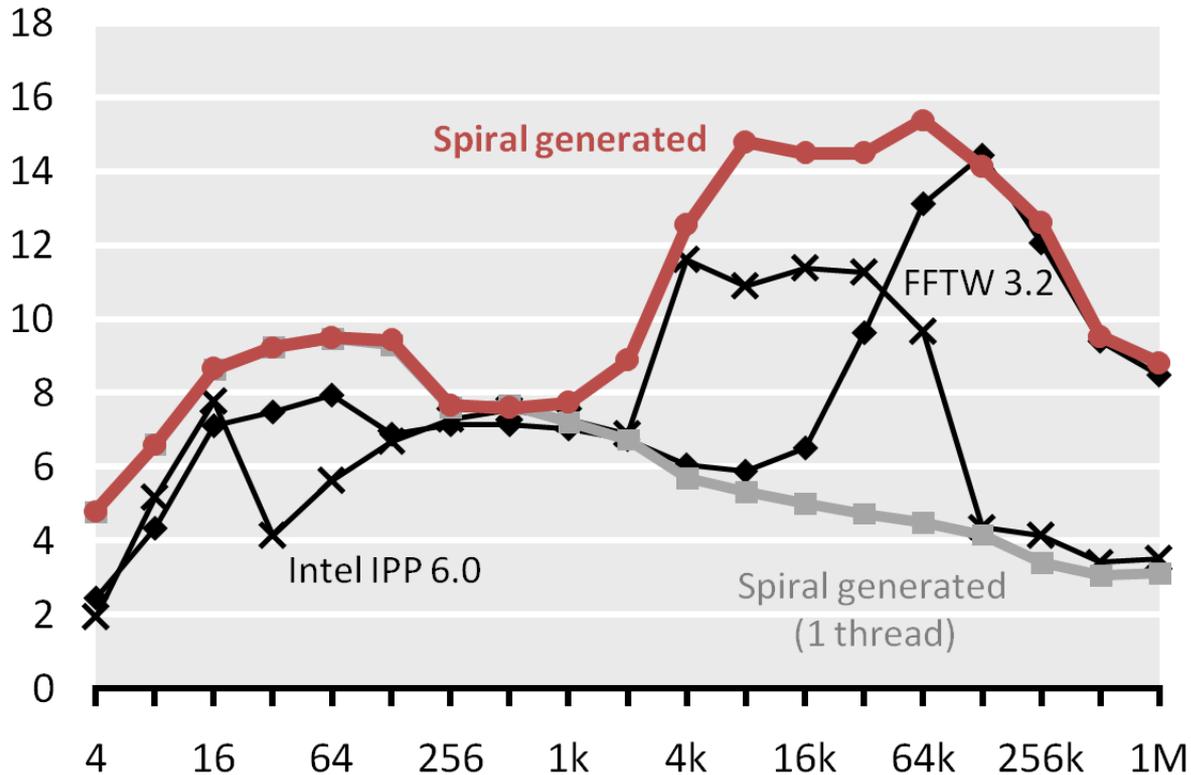
Optimized implementation
(performance/power/energy)

Spiral frees up guru programmers

DFT on Intel Multicore

Complex DFT (Intel Core i7, 2.66 GHz, 4 cores)

Performance [Gflop/s] vs. input size



$$\text{DFT}_n \rightarrow (\text{DFT}_k \otimes I_m) T_m^n (I_k \otimes \text{DFT}_m) L_k^n$$

$$\text{DFT}_n \rightarrow P_{k/2, 2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m)$$

$$\text{RDFT}_n \rightarrow (P_{k/2, m}^\top \otimes I_2) (\text{RDFT}_{2m} \oplus (I_{k/2-1} \otimes_i D_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m)$$

$$\text{rDFT}_{2n}(u) \rightarrow L_m^{2n} (I_k \otimes_i \text{rDFT}_{2m}((i+u)/k)) (\text{rDFT}_{2k}(u) \otimes I_m)$$

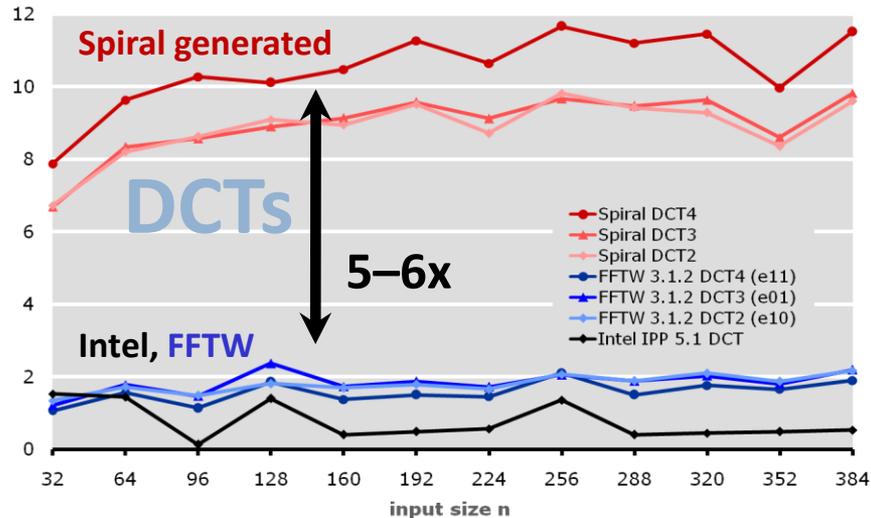


**Vectorized, threaded,
general-size, adaptive library**

Spiral Outperforms Human Programmers

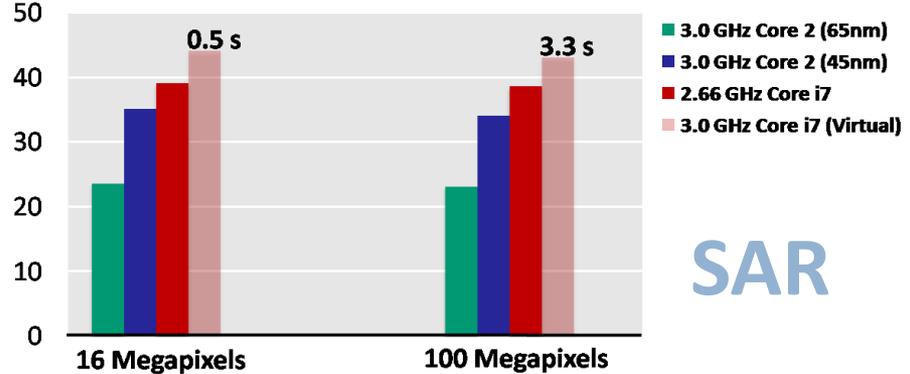
DCT on 2.66 GHz Core2 (single-precision, 4-way SSSE3)

performance [Gflop/s], opcount = 2.5 n ld n, Windows XP 32-bit, Intel C++ 10.1



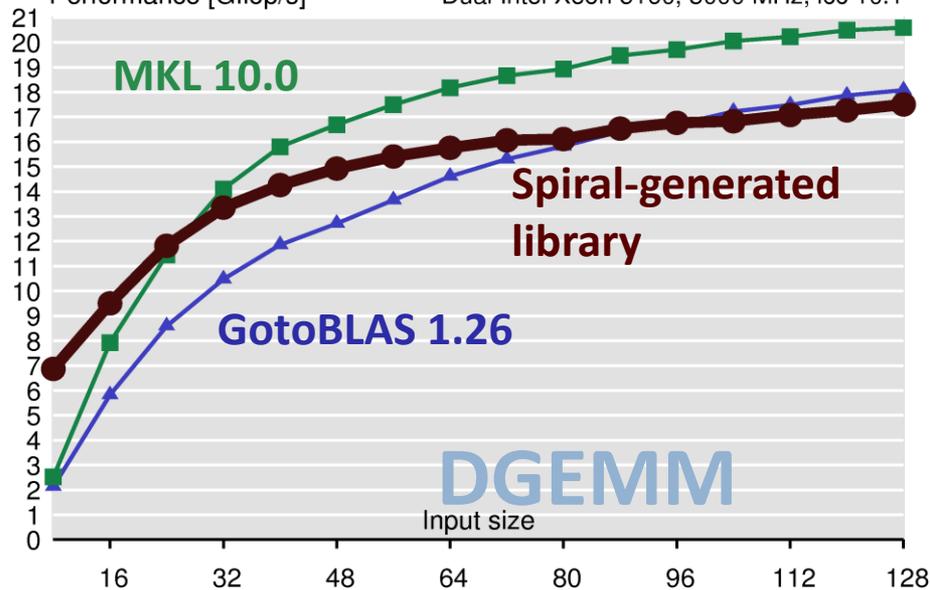
SAR Image Formation on Intel platforms

performance [Gflop/s]



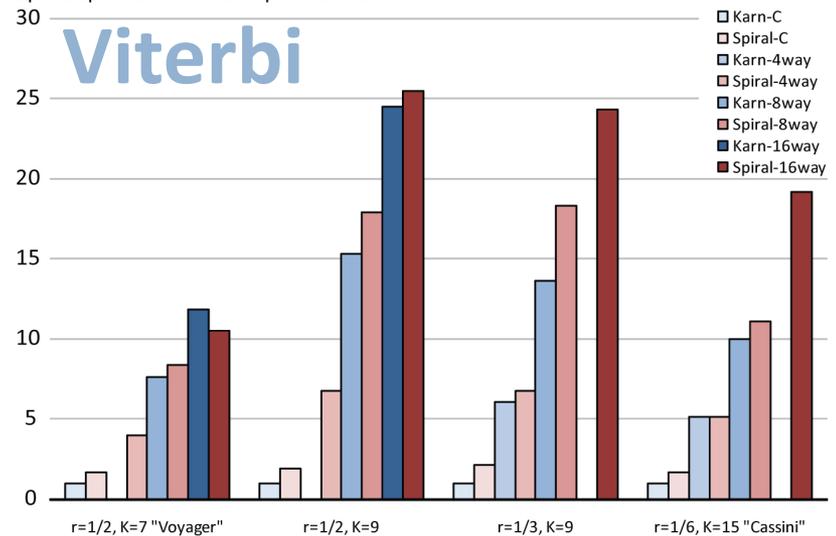
GEMM (acc. packed sq. NN), single-threaded, single precision

Performance [Gflop/s] Dual Intel Xeon 5160, 3000 MHz, icc 10.1

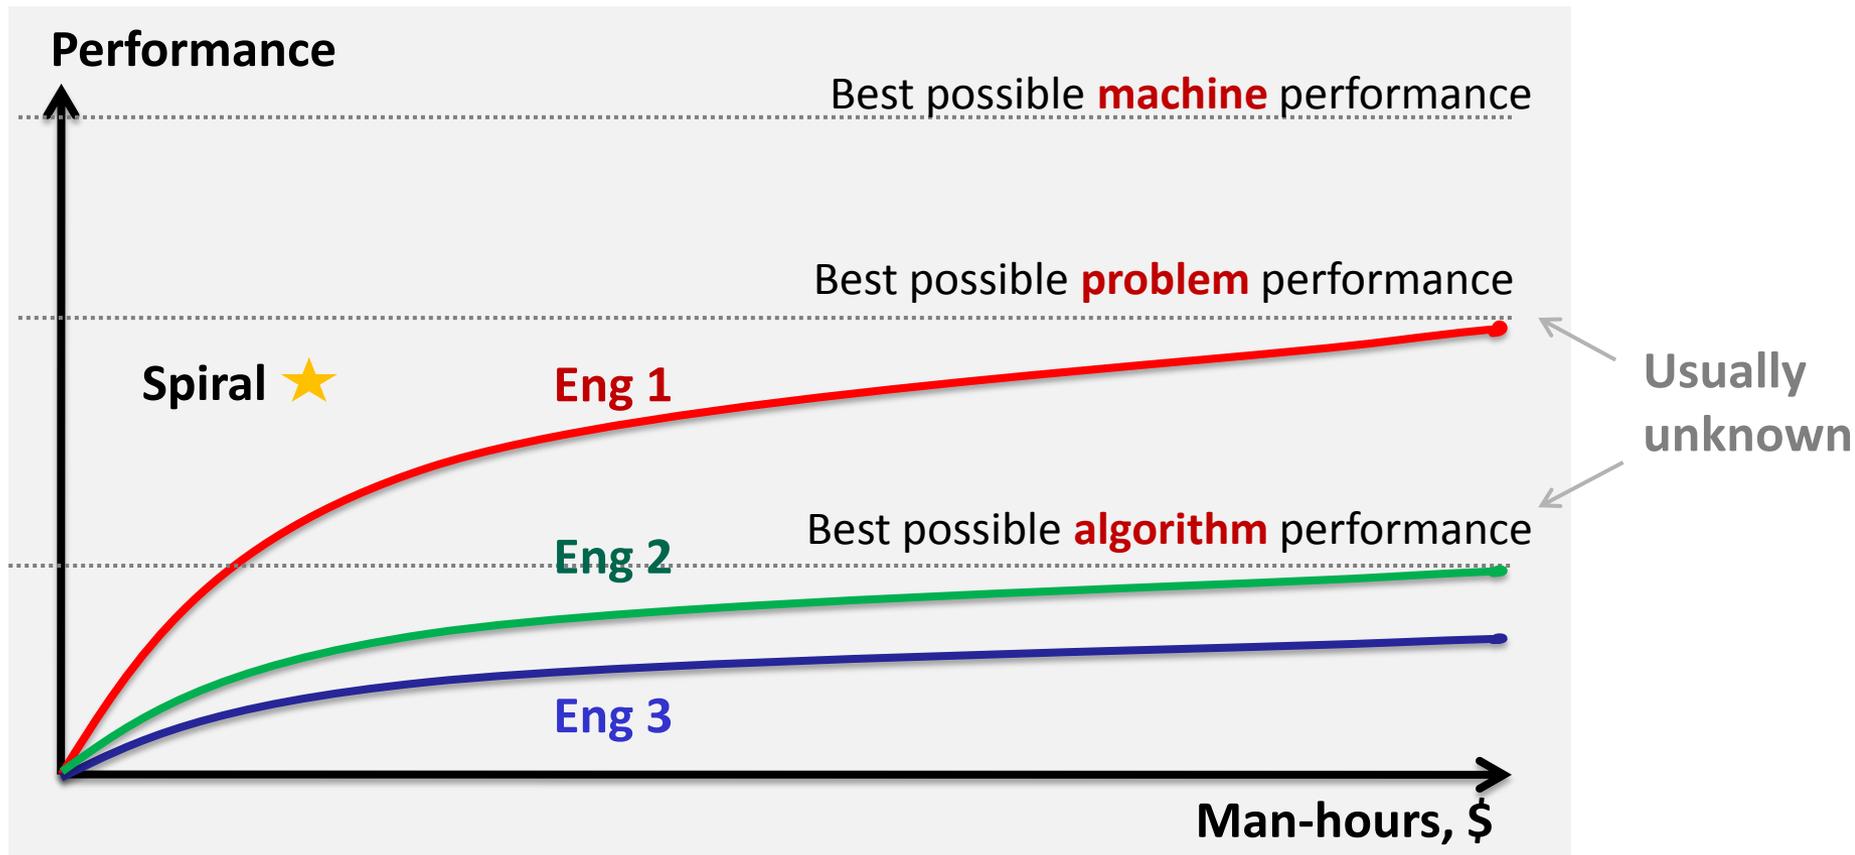


Performance Gain of Various Generated Viterbi

Speedup over Karn's C implementation



How Can Machine-Generated Code Outperform Hand-Optimized?



■ Spiral explores both

- Algorithms
- Code restructurings

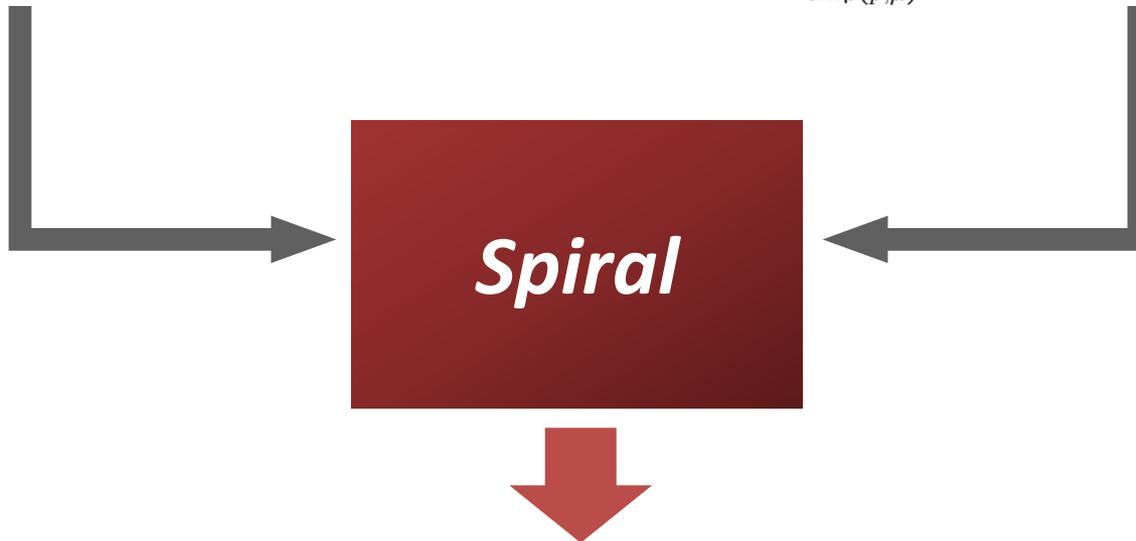
Spiral: How It Works I

Application specification
Algorithm knowledge

$$\begin{aligned} \text{DFT}_n &\rightarrow P_{k/2,2m}^\top \left(\text{DFT}_{2m} \oplus \left(I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k) \right) \right) \left(\text{RDFT}'_k \otimes I_m \right) \\ \begin{matrix} \text{rDFT}_{2n}(u) \\ \text{rDHT}_{2n}(u) \end{matrix} &\rightarrow L_m^{2n} \left(I_k \otimes_i \begin{matrix} \text{rDFT}_{2m}((i+u)/k) \\ \text{rDHT}_{2m}((i+u)/k) \end{matrix} \right) \left(\begin{matrix} \text{rDFT}_{2k}(u) \\ \text{rDHT}_{2k}(u) \end{matrix} \otimes I_m \right) \\ \text{RDFT-3}_n &\rightarrow (Q_{k/2,m}^\top \otimes I_2) (I_k \otimes_i \text{rDFT}_{2m})(i+1/2/k) (\text{RDFT-3}_k \otimes I_m) \end{aligned}$$

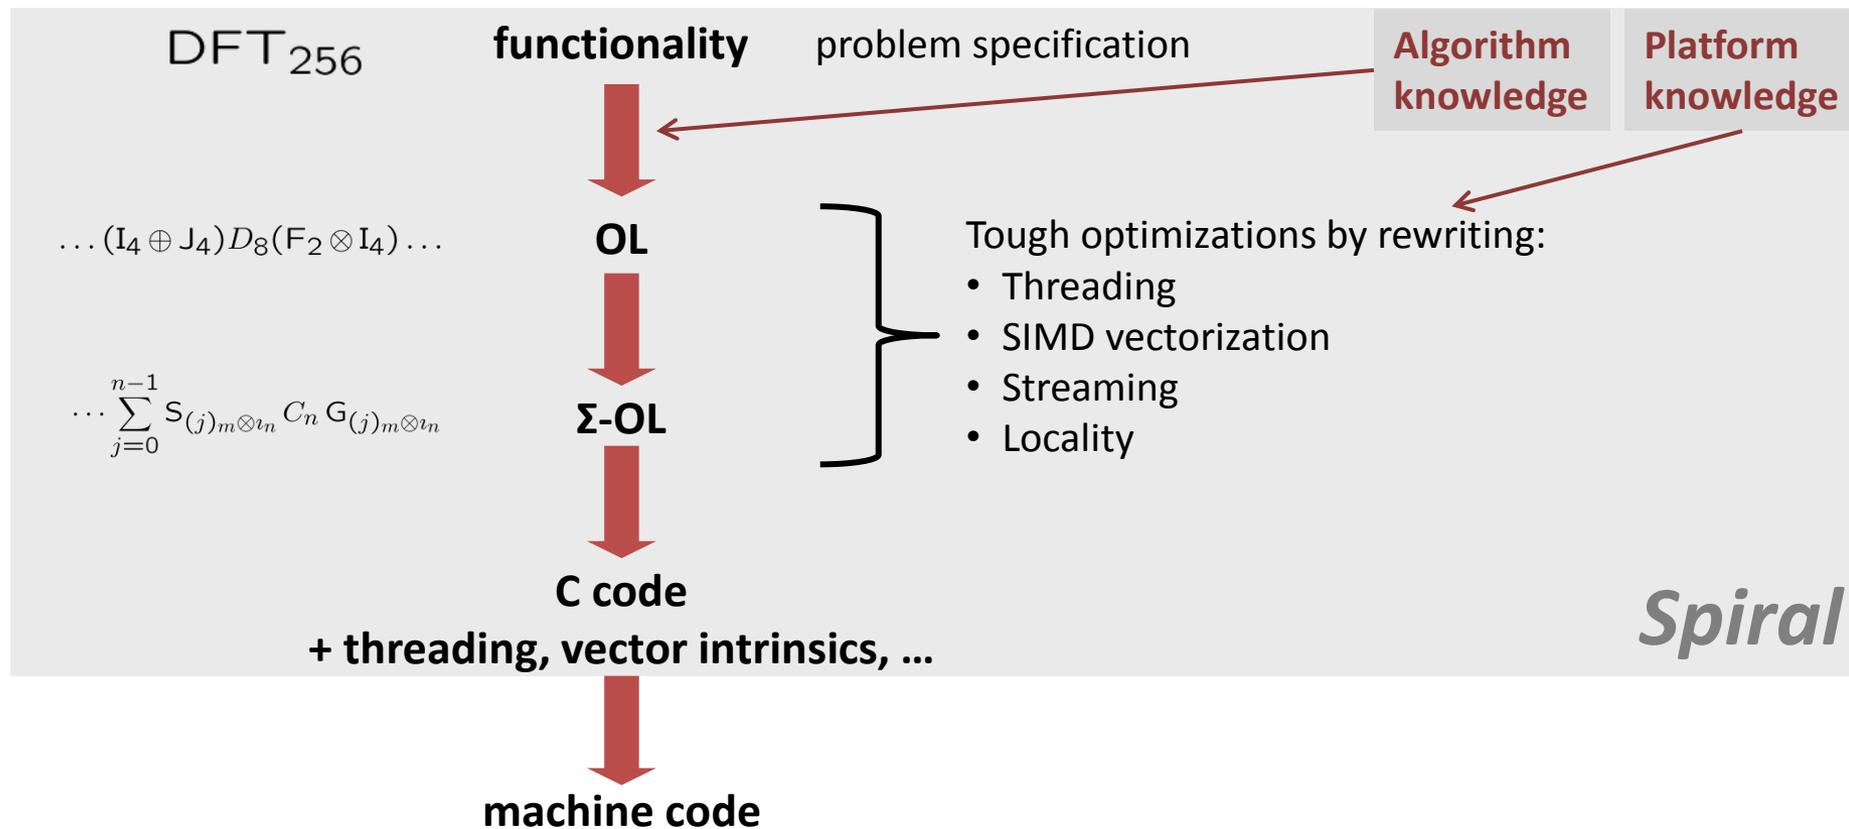
Platform description
Instruction set architecture

$$\begin{aligned} \underbrace{A_m \otimes I_n}_{\text{smp}(p,\mu)} &\rightarrow \underbrace{\left(L_m^{mp} \otimes I_{n/p} \right) \left(I_p \otimes (A_m \otimes I_{n/p}) \right) \left(L_p^{mp} \otimes I_{n/p} \right)}_{\text{smp}(p,\mu)} \\ \underbrace{I_m \otimes A_n}_{\text{smp}(p,\mu)} &\rightarrow I_p \otimes_{\parallel} \left(I_{m/p} \otimes A_n \right) \\ \underbrace{(P \otimes I_n)}_{\text{smp}(p,\mu)} &\rightarrow (P \otimes I_{n/\mu}) \overline{\otimes} I_\mu \end{aligned}$$



Optimized implementation
(performance/power/energy)

Spiral: How It Works II



- **Optimization at the high level of abstraction:**
Overcomes compiler limitations
- **Complete automation**

Example: Viterbi decoder

Select convolutional code

Select a preset code or customize parameters

- custom*
- Voyager
- NASA-DSN
- CCSDS/NASA-GSFC
- WiMax
- CDMA IS-95A
- LTE (3GPP - Long Term Evolution)
- UWB (802.15)
- CDMA 2000
- Cassini
- Mars Pathfinder & Stereo

rate 1 /
K
polynomials

code rate [\(?\)](#)
constraint length [\(?\)](#)
polynomials for the
code in decimal notation
[\(?\)](#)

Select implementation options

frame length

unpadded frame length in bits [\(?\)](#)

Vectorization level

type of code [\(?\)](#)

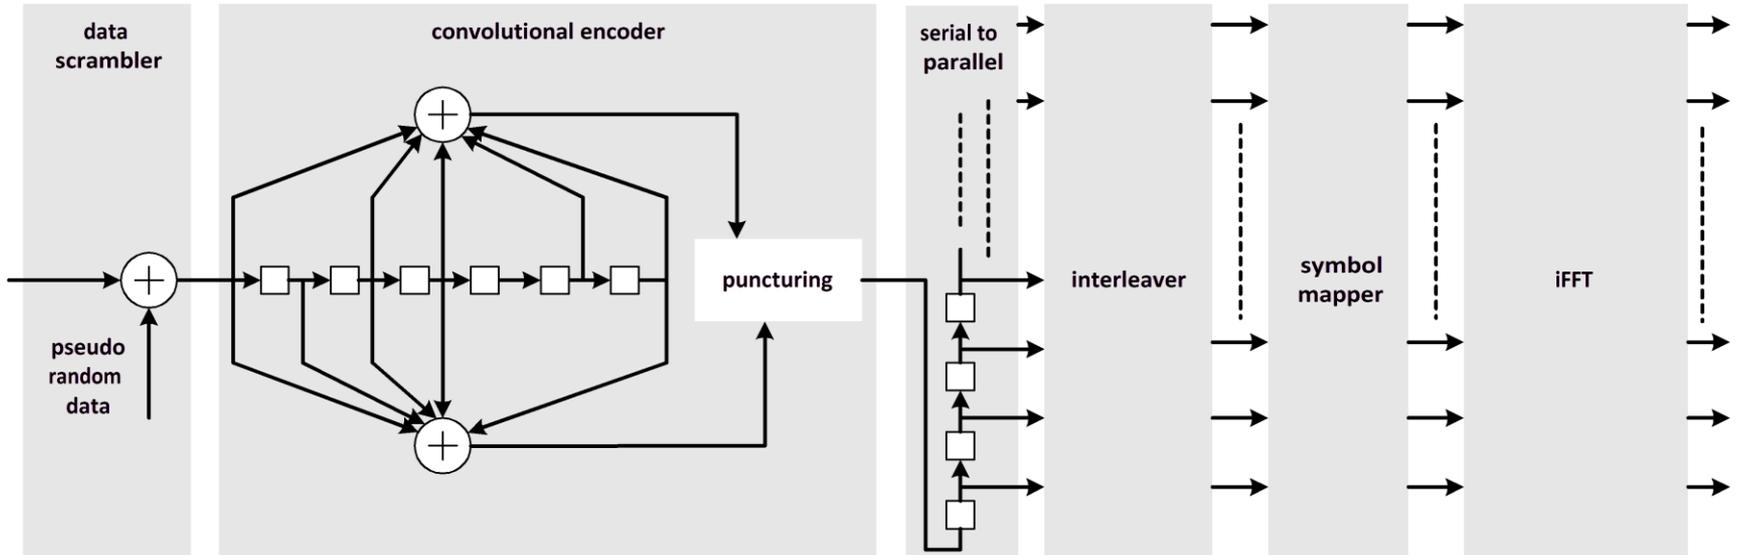
Generate Code

Reset



“Click”: Push-button code generation

“Teaching” Spiral SDR: Algorithm Knowledge



$$\text{WiFiTx}_{k,r,m} \rightarrow \left(\mathbf{I}_k \otimes (\text{CP-iDFT}_{64} \circ \text{SymbM}_m \circ \text{Interl}_m) \right) \circ \text{ConvEnc}_{r,\ell} \circ \text{Scramb}_{i_0,\ell}$$

$$\text{Scramb}_{i_0,\ell} \rightarrow (\mathbf{I}_\ell \otimes_i C_i) := \bigoplus_{i=i_0}^{i_0+\ell-1} C_i$$

$$\text{ConvEnc}_{r,\ell} \rightarrow \left(\mathbf{T}_{r,\ell+6} L_{\ell+6}^{2(\ell+6)} \begin{bmatrix} \mathbf{C}_{0,\ell} \\ \mathbf{C}_{1,\ell} \end{bmatrix} \right)$$

$$\text{Interl}_m \rightarrow \left((\mathbf{I}_{16} \otimes_i (\mathbf{I}_3 \otimes \mathbf{W}_i^m)) L_{16}^{48m} \right)$$

$$\text{SymbM}_m \rightarrow G \circ [\mathbf{I}_{48} \otimes \mathbf{S}_m]$$

$$\text{CP-iDFT}_N \rightarrow \left(\frac{1}{\sqrt{N}} P_N \text{iDFT}_N \right)$$

plus

**operator definitions
and additional rules**

“Teaching” Spiral SDR: Platform Knowledge

■ Rewriting rules

- Parallelization
- SIMD vectorization
- Cross block optimizations
- SIMD ISA specific simplification

■ Backend extensions

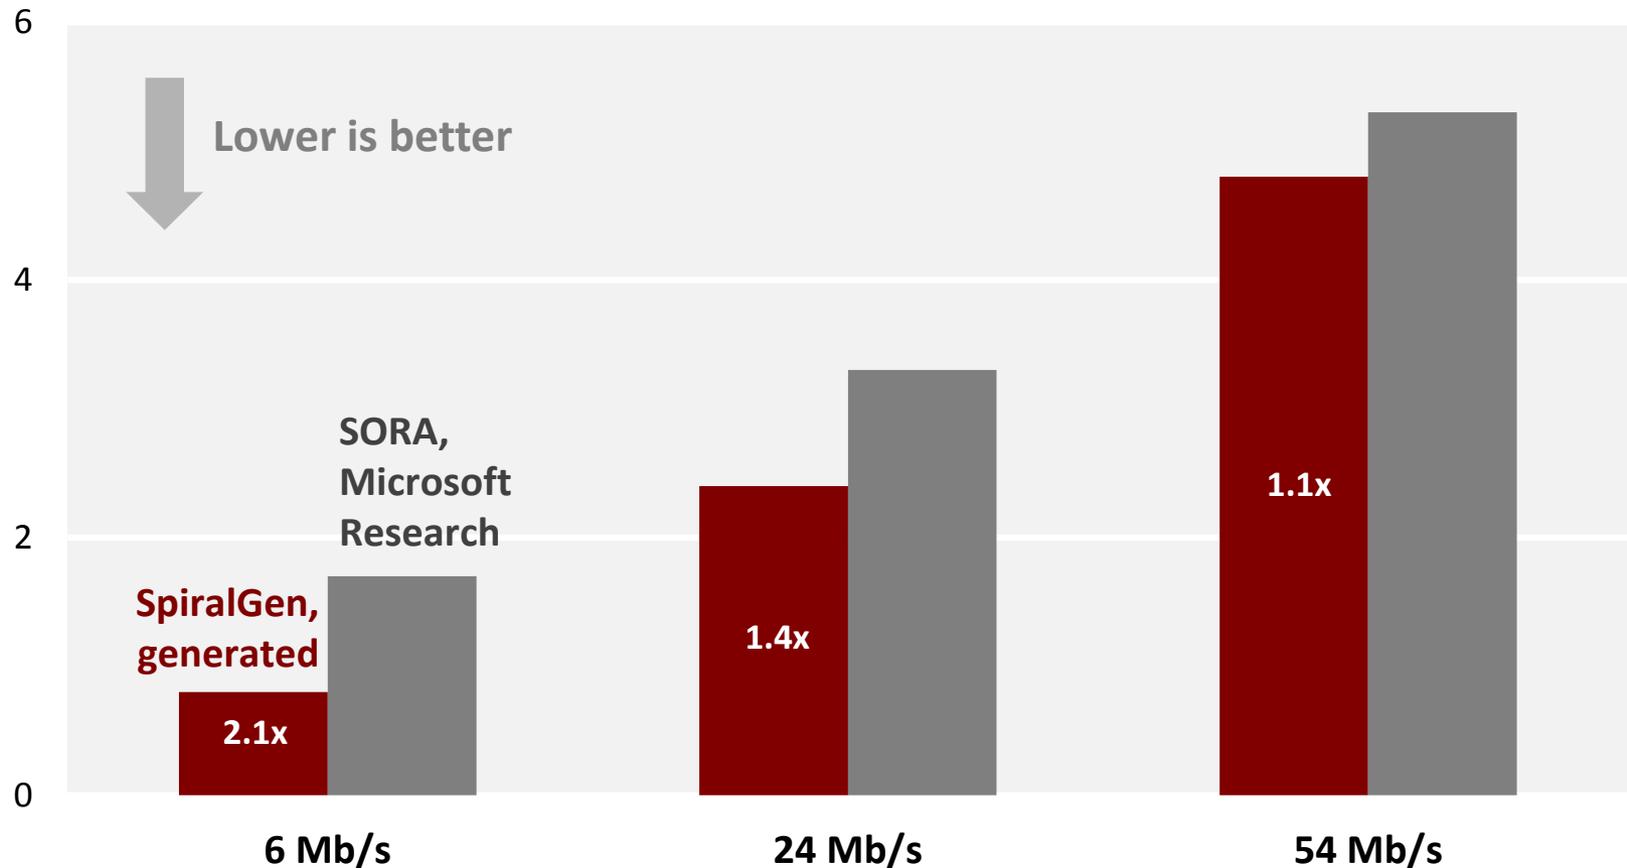
- **New data types:** bit, integer (8, 16, 32)
- **Vectorization modes:** 128-way, 16-way, 8-way, 4-way (ISAs)
- **ISA bridges:** data types conversion in vector registers
- Support for finite field arithmetic

What We Did in a Nutshell

- **WiFi 802.11.g (OFDM), 8 data rate modes**
 - Expressed in Spiral's OL
 - Extended OL for mixed data type formulas
 - Extended OL for mixed vector-length formulas
- **Created optimized implementations w/ Spiral**
 - Threading
 - SIMD vectorization, with **mixed data types / vector length**
 - Auto-tuning using best algorithm search
- **Benchmarked on 3 platforms**
 - Intel Atom
 - Intel Core 2
 - Intel Core i7

802.11b Receiver: Generated vs. Hand-tuned

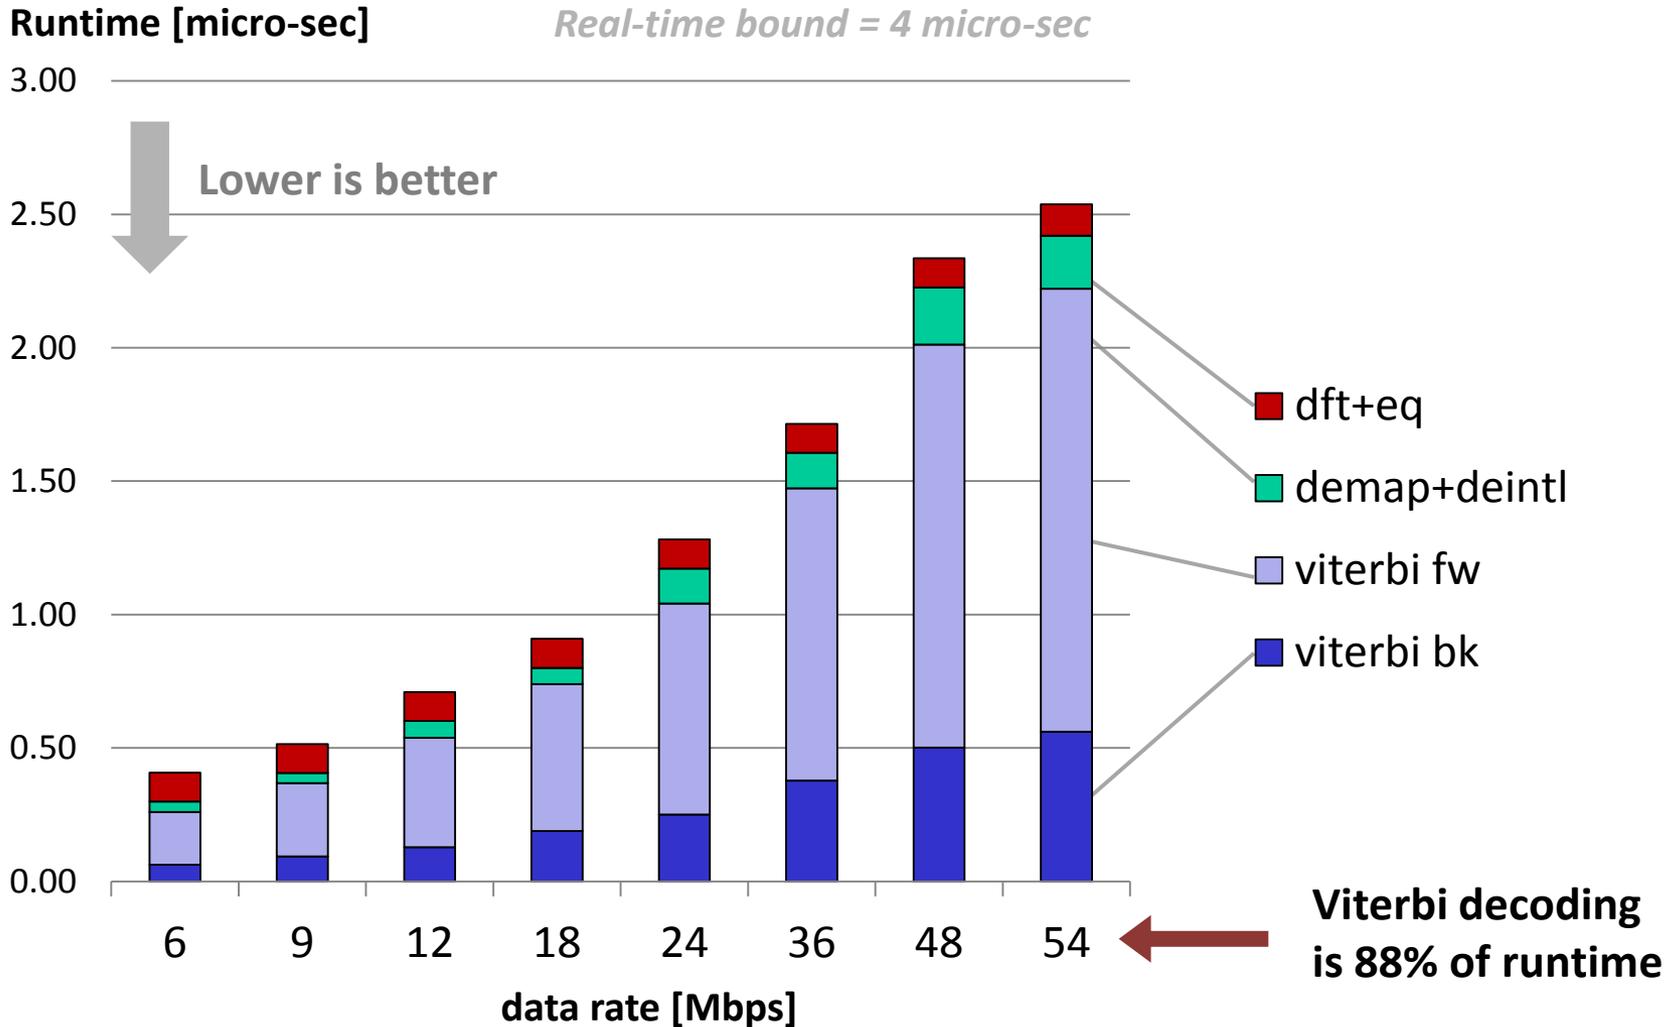
Runtime per symbol [ms]



2.1x more energy efficient
(Mflops/watt)

Compute Time Distribution

Auto-generated receiver on Core i7

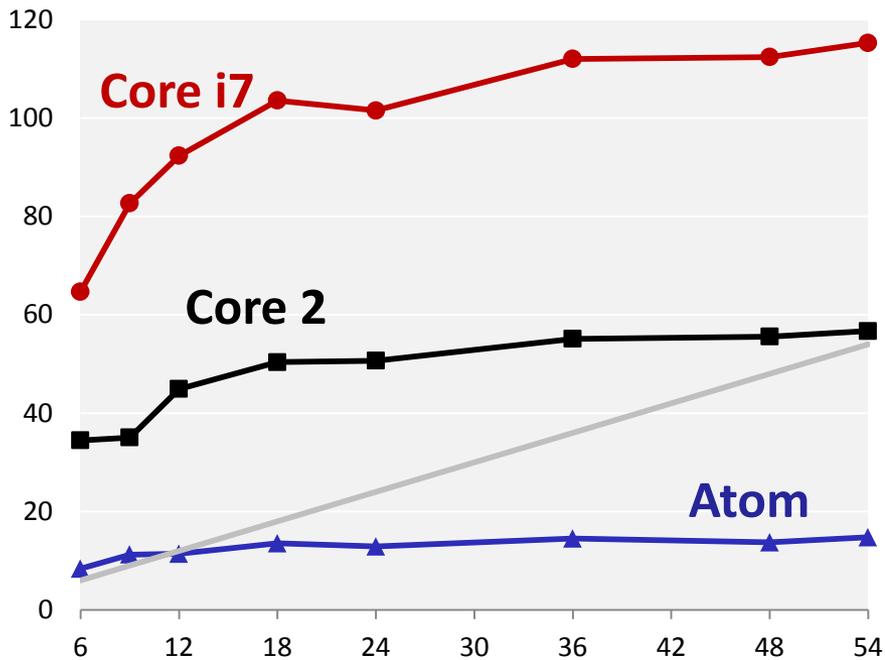


Compute-bound Bandwidth

Auto-generated WiFi receiver/transmitter (parallel)

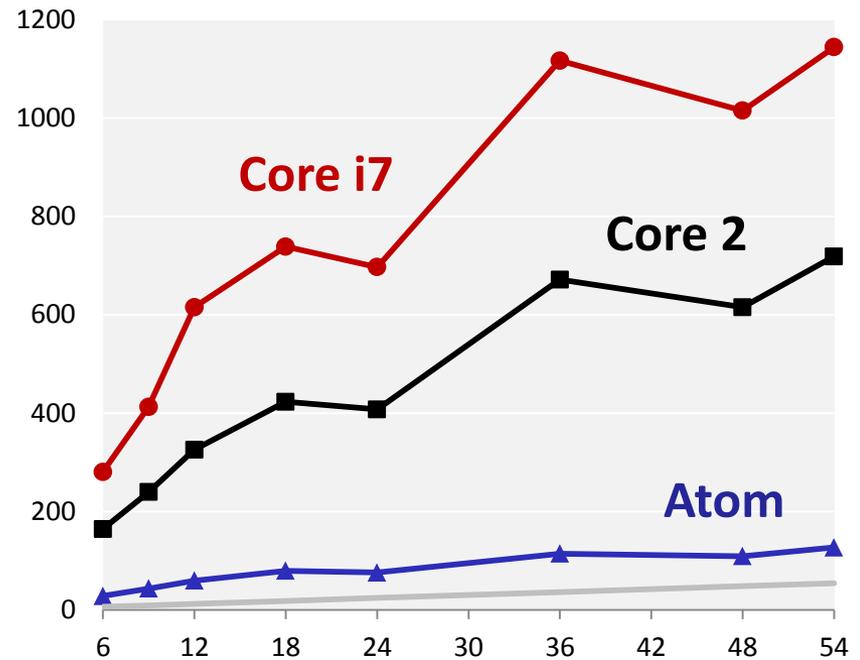
Receiver

Throughput [Mbit/s]



Transmitter

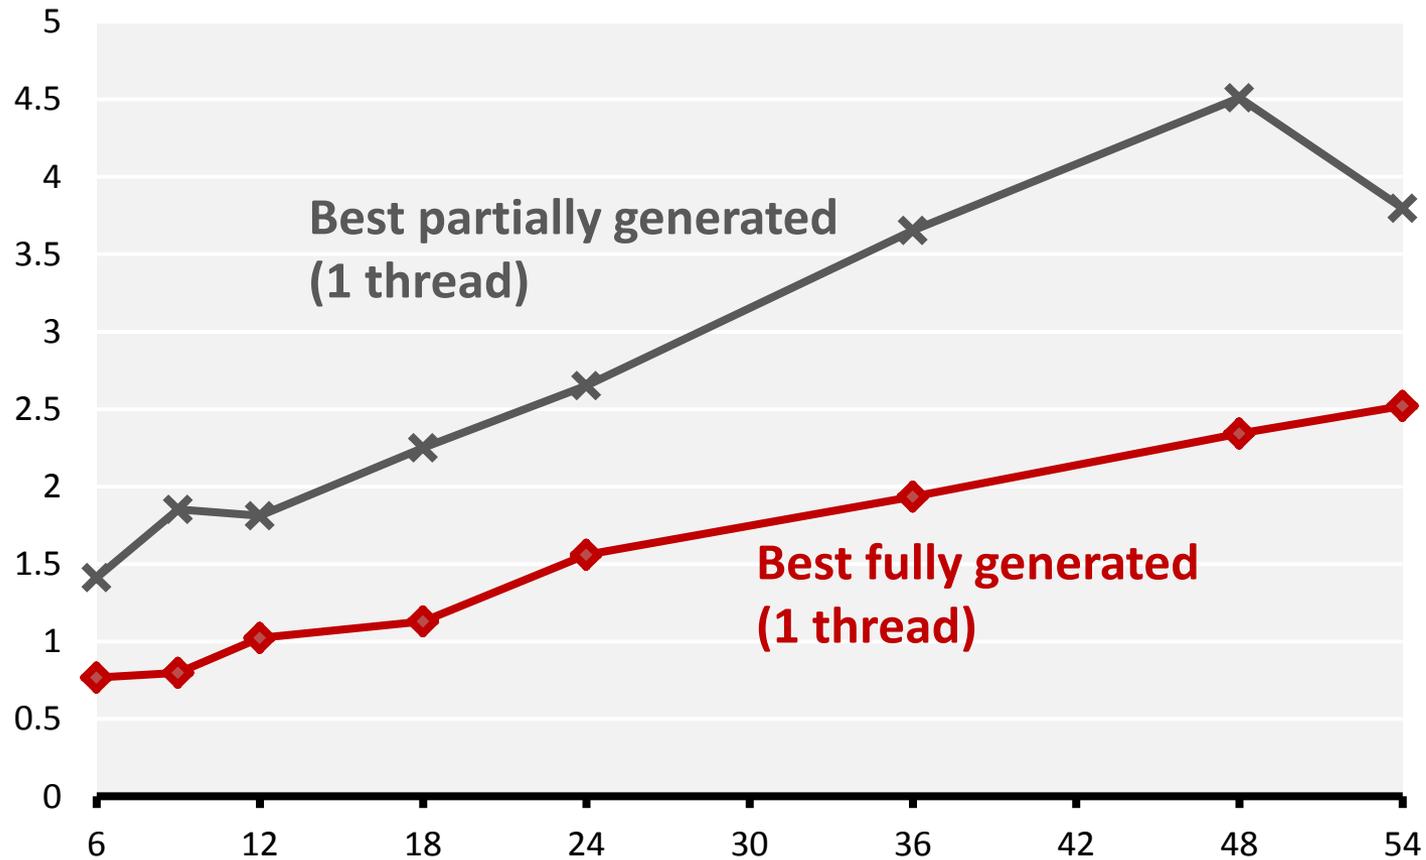
Throughput [Mbit/s]



Ability to Cross-Block Optimize

Spiral Transmitter on Atom

Run time per symbol [micro seconds] vs. Data rate [Mbit/s]



Conclusions

- The generated code is *very fast* (often faster than any human written code)
- Spiral supports *different types of architectures* including vector, multicore and distributed platforms
- Spiral enables *very fast transition to new architectures*

