





High Performance: How DSLs Can Help

Markus Püschel

Computer Science
ETH zürich

Computing



Science simulations

Audio, image, Video processing

Signal processing, communication, control

Security

Machine learning, data analytics

Optimization

***Highest performance
is often crucial***

How Do We Get Fast Code?

Algorithms

Choose cheap algorithm

Software

Implement in C/C++

Compilers

Choose good compiler and flags

Microarchitecture


Runs very fast

How well does this work?

Example: Discrete Fourier Transform

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


Performance [Gflop/s]



Example: Discrete Fourier Transform

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

Performance [Gflop/s]




Vendor compiler, best flags

Example: Discrete Fourier Transform

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


Performance [Gflop/s]



Vendor compiler, best flags
Roughly same operations count

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

Performance [Gflop/s]




Compiler doesn't do the job

Doing by hand = restructure algorithm for locality & parallelism,
handle choices, choose proper code style, use vector intrinsics,
= nightmare

Model predictive control	Singular-value decomposition
Eigenvalues	Mean shift algorithm for segmentation
LU factorization	Stencil computations
Optimal binary search organization	Displacement based algorithms
Image color conversions	Motion estimation
Image geometry transformations	Multiresolution classifier
Enclosing ball of points	Kalman filter
Metropolis algorithm, Monte Carlo	Object detection
Seam carving	IIR filters
SURF feature detection	Arithmetic for large numbers
Submodular function optimization	Optimal binary search organization
Graph cuts, Edmond-Karps Algorithm	Software defined radio
Gaussian filter	Shortest path problem
Black Scholes option pricing	Feature set for biomedical imaging
Disparity map refinement	Biometrics identification

Same for (almost) all computational problems:
Straightforward code is highly suboptimal


Current



C code is a singularity:

- Compiler has no access to high level information
- No structural optimization
- No evaluation of choices

Future



Challenge: conquer the high abstraction level for ***more/complete automation***

DSLs!

Example: Spiral

Computer Generation of Fast DFTs

www.spiral.net

Recursive algorithms expressed as rules in mathematical, internal DSL

$$\begin{aligned}\text{DFT}_n &\rightarrow (\text{DFT}_k \otimes I_m) T_m^n (I_k \otimes \text{DFT}_m) L_k^n \\ \text{DFT}_n &\rightarrow P_{k/2,2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m) \\ \text{RDFT}_n &\rightarrow (P_{k/2,m}^\top \otimes I_2) (\text{RDFT}_{2m} \oplus (I_{k/2-1} \otimes_i D_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m) \\ \text{rDFT}_{2n}(u) &\rightarrow L_m^{2n} (I_k \otimes_i \text{rDFT}_{2m}((i+u)/k)) (\text{rDFT}_{2k}(u) \otimes I_m)\end{aligned}$$

Recursive combination yields many choices

Example: Spiral

Transform

DFT₈



Decomposition rules

Algorithm
(DSL 1)

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



parallelization
vectorization

Algorithm
(DSL 2)

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



C Program
+ ext.


```
void sub(double *y, double *x) {  
    double f0, f1, f2, f3, f4, f7, f8, f10, f11;  
    ...  
    t282 = _mm_addsub_ps(t268, U247);  
    t283 = _mm_add_ps(t282, _mm_addsub_ps(U247, _mm_shuffle_ps(t275,  
    t284 = _mm_add_ps(t282, _mm_addsub_ps(U247, _mm_sub_ps(_mm_setze  
    s217 = _mm_addsub_ps(t270, U247);  
    s219 = _mm_shuffle_ps(t278, t280, _MM_SHUFFLE(1, 0, 1, 0));  
    s220 = _mm_shuffle_ps(t278, t280, _MM_SHUFFLE(3, 2, 3, 2));  
    s221 = _mm_shuffle_ps(t283, t285, _MM_SHUFFLE(1, 0, 1, 0));  
    ...  
    < many more lines>
```



locality
optimization

+ *Search or
Learning for
Choices*

Example: Delite



Generating Fast Database Code with DBLAB

Maps query/transaction workloads to embedded Scala DSL.

DSL compiler with a rich set of domain-specific code transformers (data layout transformations, data structure specialization, index introduction, materialization decisions, ...)

Uses code transformations on multiple abstraction levels. Successive lowering phases.

Generates fast C code

