

GLAF: A Visual Programming and Auto-Tuning Framework for Parallel Computing

Student: Konstantinos Krommydas Collaborator: Dr. Ruchira Sasanka (Intel) Advisor: Dr. Wu-chun Feng





Motivation

High-performance computing is crucial in a broad range of scientific domains:

• Engineering, math, physics, biology, .





Parallel programming

revolution has made high-performance computing **accessible** to the broader masses





Challenges

Many different computing platforms

Many different programming languages

Many different architectures

Many different optimization strategies

Domain experts should not (need to) know all these details

but rather focus on their science





Challenges

Domain experts need to collaborate with computer scientists



- Communication overhead & errors
- Need to exchange domainspecific/programming knowledge

Invent the Future

Innovation slow-down

Limited access to parallel computing



Contributions

- Realize a programming abstraction & development framework for domain experts to provide a balance between performance and programmability
 - i.e., obtain fast performance from algorithms that have been programmed easily

Desired features

intuitive, familiar, minimalistic syntax

data-visual and interactive

auto-parallelizable, optimizable, tunable

able to integrate with existing legacy code



*GLAF

*Grid-based Language and Auto-tuning Framework



Programming Using GLAF: a Simple Example

	(integer)								
	0	1	2	3	4	5	б		
0	0	0	0	0	0	0	0		0
1	0	24	50	78	108	140	174		1
2	0	96	200	312	432	560	696		2
3	0	216	450	702	972	1260	1566		3
4	0	384	800	1248	1728	2240	2784		4
5	0	600	1250	1950	2700	3500	4350		5
6	0	864	1800	2808	3888	5040	6264		6

scaled_output

	(integer)							
	0	1	2	3	4	5	6	
0	0	0	0	0	0	0	0	
1	23	24	25	26	27	28	29	
2	46	48	50	52	54	56	58	
3	69	72	75	78	81	84	87	
4	92	96	100	104	108	112	116	
5	115	120	125	130	135	140	145	
6	138	144	150	156	162	168	174	

input

Parallelism: row:7 col:7 Insert New Step Duplicate Step

	(Integer)							
	0	1	2	3	4	5	6	
0	0	0	0	0	0	0	0	
1	0	1	2	3	4	5	6	
2	0	2	4	6	8	10	12	
3	0	3	6	9	12	15	18	
4	0	4	8	12	16	20	24	
5	0	5	10	15	20	25	30	
6	0	6	12	18	24	30	36	

Delete Step

Main Menu

scaling_data

Index Variables · [row co]]								
	Index Variables : [IOW, COI]							
Index Range:	foreach row col							
Condition:								
Formula:	<pre>scaled_output[row,col] = input[row,col] * scaling_data[row,col]</pre>	Delete Formula						
	+ - * / < > < = != <- OR AND NOT () f_{new} f f_{lib} 123 abc end Delete							



bch Invent the Future



- GLAF variables are based on the concept of grids:
 - familiar abstraction (e.g., images, matrices, spreadsheets)
 - regular format that facilitates code generation, optimizations, and parallelism detection
- Grid-based programming puts the focus on the relation, rather than the implementation details
- Example:
 - Scalar variable: 0D grid
 - 1D array: one-dimensional grid













College	e A Co	llege B	College C
Dept	A Dept. I	B Dept. C	Dept. D
	Males	Fer	males
ę	string v	string v	integer v
	first name 👎	last name 👎	age 두
	stud	lents	





GLAF Infrastructure



- GLAF Programming GUI
- Data visualization
- Fortran/C/JS/OpenCL code generation
- Auto-parallelization

roiniaTech

Invent the Future

 Compilation/auto-tuning script generation



Web/Cloud

- Auto-tuning
- Execution

- Data storage
 - 12



Code Generation

```
int ft calcPointCharge(int *ft n atoms, float *ft sum Fs, int ft curr surf pt, struct TYP surface pts typvar surface pts, struct
TYP_surface_pts typvar_atoms, float ft_Ke) {
       int fun_curr_surf_pt;
       float fun Ke;
       int ft_ReturnValue;
       int ft row;
       int ft end0;
       int atoms q=0, atoms x=1, atoms y=2, atoms z=3;
       int surface pts q=0, surface pts x=1, surface pts y=2, surface pts z=3;
       float ft r2;
       fun curr_surf_pt = ft_curr_surf_pt;
       fun Ke = ft Ke;
       // Loop through all atoms vs single surf pt
       ft end0 = 4-1;
       float tmp_sum_Fs=ft_sum_Fs[fun_curr_surf_pt];
       #pragma omp parallel for collapse(1) private(ft_r2) reduction(+: tmp_sum_Fs)
       for (ft row = 0; ft row <= ft end0; ft row += 1) {</pre>
               // Calculating the distance between a surface point and each atom
               ft r2 = ft calcDistance(typvar surface pts,fun curr surf pt,typvar atoms,ft row);
               // Add current pairs charge to total
               tmp sum Fs = tmp sum Fs + ft calcPairCharge(fun Ke,ft r2,typvar surface pts,fun curr surf pt,typvar atoms,ft row);
       }
       ft sum Fs[fun curr surf pt]=tmp sum Fs;
```



}



Code Generation

```
INTEGER FUNCTION ft calcPointCharge(ft n atoms,ft sum Fs,ft curr surf pt,typvar surface pts,typvar atoms,ft Ke)
       INTEGER, DIMENSION(:) :: ft n atoms
       REAL, DIMENSION(:) :: ft sum Fs
       INTEGER :: ft curr surf pt
       TYPE (TYP surface pts) typvar surface pts
       TYPE (TYP_atoms) typvar_atoms
       REAL :: ft Ke
       INTEGER :: fun_curr_surf_pt
       REAL :: fun Ke
       INTEGER ::ft ReturnValue
       INTEGER :: ft row
       INTEGER :: ft end0
       INTEGER :: atoms_q=0, atoms_x=1, atoms_y=2, atoms_z=3
       INTEGER :: surface_pts_q=0, surface_pts_x=1, surface_pts_y=2, surface_pts_z=3
       REAL ft r2
       fun curr surf pt = ft curr surf pt
       fun Ke = ft Ke
       !Loop through all atoms vs single surf pt
       ft end0 = 4-1
       !$OMP PARALLEL DO COLLAPSE(1) &
       !$OMP REDUCTION(+: ft_sum_Fs) &
       !$OMP PRIVATE(ft r2)
       D0 ft row = 0, ft end0, 1
               ! Calculating the distance between a surface point and each atom
               ft_r2 = ft_calcDistance(typvar_surface_pts,fun_curr_surf_pt,typvar_atoms,ft_row)
               ! Add current pairs charge to total
               ft sum Fs(fun curr surf pt + 1) = ft sum Fs(fun curr surf pt + 1) + ft calcPairCharge(fun Ke,ft r2,typvar surface pts
,fun_curr_surf_pt,typvar_atoms,ft_row)
       END DO
       *$0MP END PARALLEL DO
       ft calcPointCharge = ft ReturnValue
       RETURN
END FUNCTION
```





Auto-Tuning

Selects the *languages* in which to autogenerate code

Selects optimizations for *each* combination of language and code "starting point"

Invent the Future

```
Target Platform:
OCPU
MIC
   Gen Graphics
Target Languages:
  Fortran
   С
  OpenCL
Basic Auto-Tuning Options:
   Serial version
Parallel version (tool-generated)
Parallel version (compiler-generated)
Extra Auto-Tuning Options:
Data layout transformations (SoA/AoS)
Loop collapse transformations
Loop interchange transformations
```

Create Source

Create Binaries

Generate auto-tunescript

Auto-tune and time

Generates *platform-specific* binaries & optimizations

Selects one or more code "starting points" for **each** language



Auto-Tuning

Target Platform:

- CPU
- MIC
- Gen Graphics

Target Languages:

- Fortran
- 🗆 C
- OpenCL

Basic Auto-Tuning Options:

- Serial version
- Parallel version (tool-generated)
- Parallel version (compiler-generated)

Extra Auto-Tuning Options:

- Data layout transformations (SoA/AoS)
- Loop collapse transformations
- Loop interchange transformations

Create Source

Create Binaries

Generate auto-tunescript

Auto-tune and time

```
struct TYP surface pts {
  float *dim0;
  float *dim1:
  float *dim2:
  float *dim3;
};
// Surface points outside the biomolecule
struct TYP surface pts typvar surface pts;
// Allocation of each array within the struct (SoA)
typvar surface pts.dim0=(float *)malloc(sizeof(float )*3);
// Value assignment
typvar surface pts.dim0[ft col] = ft col * 3;
struct TYP surface pts {
  float dim0;
  float dim1;
  float dim2;
  float dim3;
};
// Surface points outside the biomolecule
struct TYP surface pts *tupvar surface pts;
// Initialization of array of structures (AoS)
typvar surface pts = (struct TYP surface pts *)malloc(
                     sizeof(struct TYP surface pts)*3);
// Value assignment
```

```
typvar_surface_pts[ft_col].dim0 = ft_col * 3;
```





Visualization

- Data visualization facilitates:
 - understanding the algorithm being developed
 - revealing bugs at an early stage

	R				G		В	
	P	0	1	2	3	4	5	6
	0	230	52	52	52	52	52	52
"Chour Data"	1	72	230	74	75	76	77	78
Show Data	2	92	94	230	98	100	102	104
	3	112	115	118	230	124	127	130
	4	132	136	140	144	230	152	156
	5	152	157	162	167	172	230	182
	6	172	178	184	190	196	202	230





Visualization

- Data visualization facilitates:
 - understanding the algorithm being developed
 - revealing bugs at an early stage







Visualization

- Data visualization facilitates:
 - understanding the algorithm being developed
 - revealing bugs at an early stage







Results: 3D finite difference algorithm





VirginiaTech

Results: N-body algorithm







Related Work



- + Includes loop-level parallelism & data-level parallelism (vectorization), and potential optimizations
- Requires certain programming knowledge
- (Often) conservative nature



Problem	Domain-
Solving	Specific
Environments	Languages

Domain-specific:

computational biology, physics, dense linear algebra, Fast-fourier transforms,

- + High-performance auto-tuning
- Restrictive in nature
- Restrictive in terms of target language and/or platform



Future Work

- Improve tool's robustness
- Enabling more languages/extensions:
 - OpenCL, OpenACC
 - Support of distributed programming (MPI)
- Dynamic feedback/advice on parallelism issues
- Extend auto-tuning, auto-parallelization/auto-vectorization capabilities
- Implement more dwarfs and provide back-end support for common programming pitfalls in code generated for supported languages





Conclusion

Invent the Future

GLAF targets domain experts and provides a fine balance between **performance** and **programmability**

- auto-parallelization, optimization and auto-tuning
- helps avoid common programming pitfalls



"GLAF: A Visual Programming and Auto-Tuning Framework for Parallel Computing" *Krommydas, Sasanka, Feng*

