

Visualization of MPI(-IO) Datatypes

Julian M. Kunkel, Thomas Ludwig

julian.martin.kunkel@informatik.uni-hamburg.de

Scientific Computing
Department of Informatics
University of Hamburg

30.08.2011

Scope of this presentation

- Visualization of MPI datatypes and non-contiguous I/O in a post-mortem trace analyzer.
- General scheme to illustrate datatypes and I/O access.

1 Introduction

2 HDTrace

3 Visualization of datatypes

4 Summary

Derived datatypes in MPI

Motivation

- Allow to communicate non-contiguous and mixed datatypes
- File-views in MPI-2 permit non-contiguous I/O

Construction of datatypes

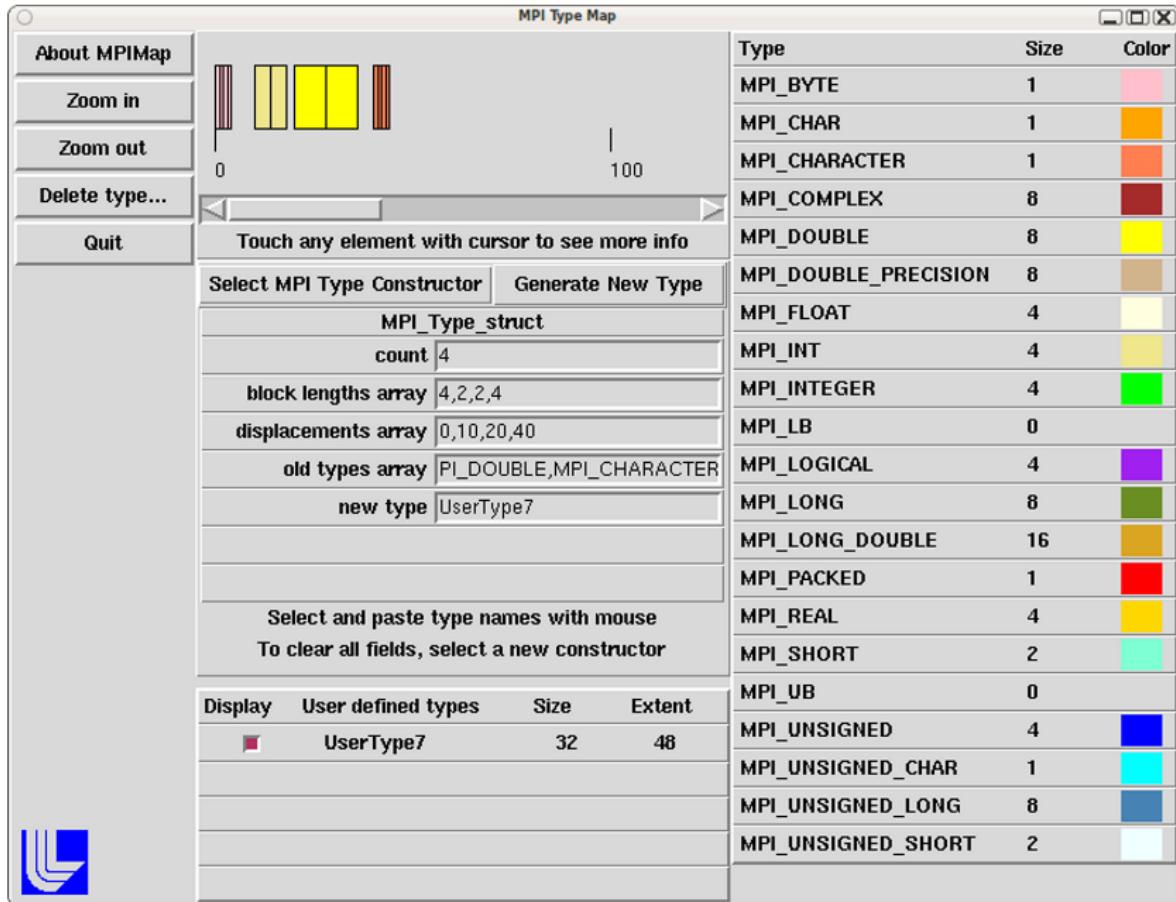
- Datatypes are created bottom up
- Use a constructor and defined datatypes ⇒ new datatype
- Constructors: contiguous, structure, vector, indexed ...

The need to analyze datatypes

- Nested datatypes may be complicated
- Debugging of datatypes
 - Check the right bytes in memory or on disk are accessed
- Reverse engineering

MPIMAP

- TCL/TK tool which visualizes the creation of datatypes
- Interactively create up to 4 datatypes
- Unrolls the accessed area on the x-axis
- Runs with a single MPI process



1 Introduction

2 HDTrace

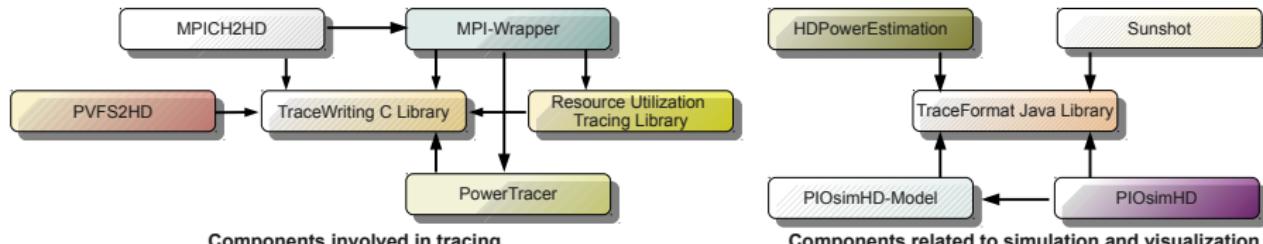
3 Visualization of datatypes

4 Summary

Overview

Research environment for tracing and simulation

- Analyze energy efficiency and I/O activity
- Modeling and simulation of systems executing MPI-IO apps
- Suitable for experiments of medium scale
- XML based trace format with user defined semantics



Components involved in tracing

Components related to simulation and visualization

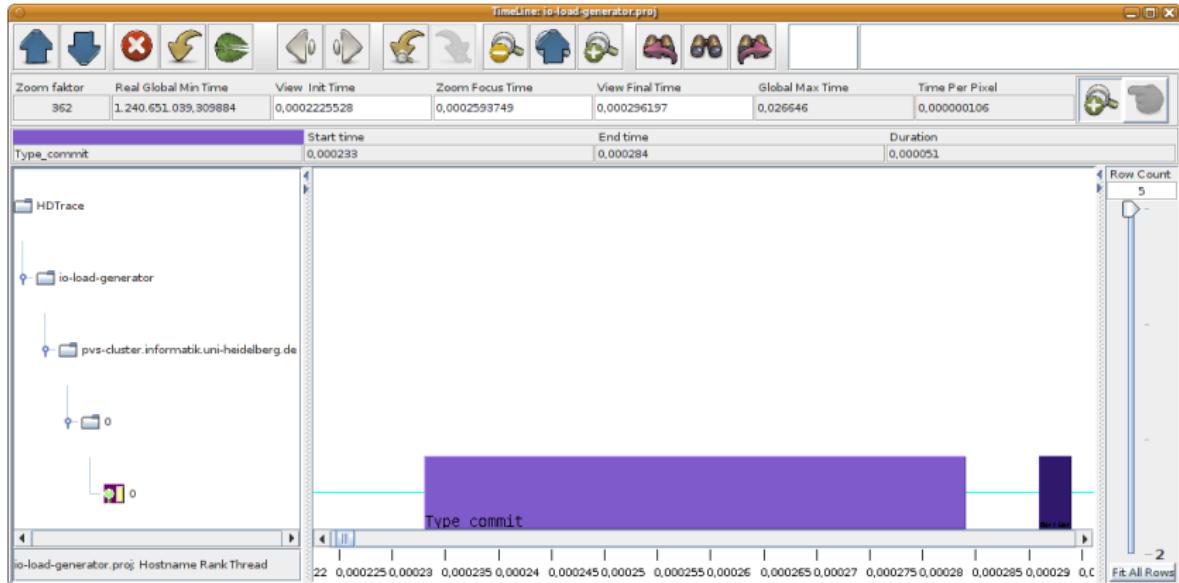
Sunshot

Overview

- Viewer for traces (and profiles)
- Written in Java, based on Jumpshot GUI

Features

- Renders numeric values (statistics) and events in timelines
- Topology maps (named) trace data to timelines
- Interaction with statistics: Histograms, derived timelines, ...
- User supplied filtering of events



1 Introduction

2 HDTrace

3 Visualization of datatypes

4 Summary

Visualization of datatypes

To start the visualization

- Select an event in the timeline
- Right click to open *TraceEntry information box*

Internal activity invoked upon startup

- Plugins scan the stored attributes of the trace entry
 - When a project is loaded plugins decide if they should be enabled
 - Depends on the topology labels (semantics)
- Memory and I/O datatypes are detected and visualized
 - Logic is implemented according to the MPI datatype constructors

TraceEntry Info Box

Trace entry	
Name	Type_commit
Start [t]	0.000233s
Duration [t]	0.000051s

Memory Datatype

STRUCT <50, 2600>			
1 x	5 B	50 x	2545 B
LB	BYTE		UB

Contained XML data:

```
<Type_commit end="1240651036.310168" time="1240651036.310117" tid="-1946157051">
</Type_commit>
```

close

Advanced topics

Nested datatypes

- Draw the (final) memory datatype
- Render child datatypes in a directed graph
- Expansion of the child datatypes is requested by the user
 - Perform a click on the datatype to expand it

Limitations

Ascending offsets in the datatype (limitation to MPI-IO as well).

Trace entry info box

Trace entry	File set view		
Name	0.022471s		
Start [t]	0.000007s		
Duration [t]			
File operation	/tmp/test-test		
file name	File datatype		
STRUCT <1066, 56500>			
1 x 4 B LB	4 x INTEGER	20 B VECTOR	1 x 960 B STRUCT

Elementary file datatype

BYTE

Contained XML data:

```
<File_set_view fid="0" time="1240651036.332355" filetid="-1946157050" representation="native" etid="1275068685" offset="0" end="1240651036.332362">
<Info value="CREATE_SET_DATAFILE_NODES">
</Info>

<Info value="striping_unit">
</Info>

</File_set_view>
```

STRUCT <1066, 56500>

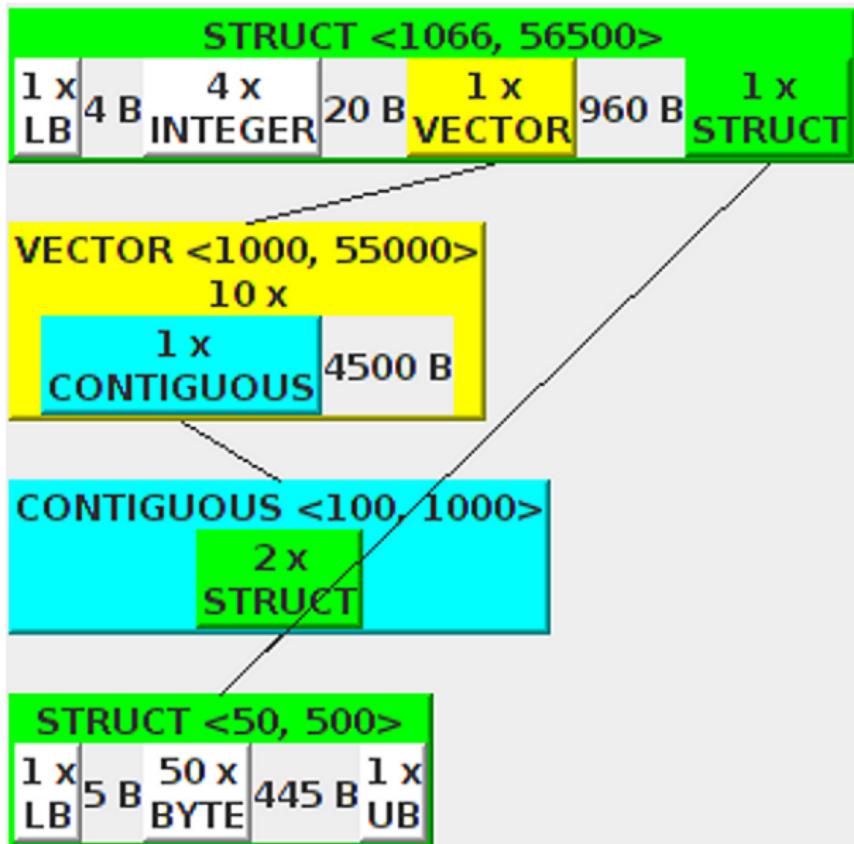
1 x	4 B	4 x	20 B	1 x	960 B	1 x
LB	INTEGER			VECTOR		STRUCT

STRUCT <50, 500>

1 x	5 B	50 x	445 B	1 x
LB	BYTE			UB

VECTOR <1000, 55000>

10 x	
1 x	4500 B
CONTIGUOUS	



Non-contiguous I/O

Visualization

- Filetype is drawn similar to memory datatypes
- Unrolling of datatype to visualize accessed regions

Datatype unrolling

- Non-contiguous I/O might hit file datatype partially
 - Write 20 bytes, but file datatype addresses 100 byte
- Three regions: partially accessed datatype, fully repeated datatype, partially accessed datatype

TraceEntry Info Box

Trace entry

Name	File_write_at
Start [t]	0.022478s
Duration [t]	0.000624s

File operation

file name	/tmp/test-test
size	501760
physical offset (after view)	501760
etype size, extent	1, 1

Memory Datatype

File datatype

STRUCT <1066, 56500>

1 x	4 B	20 B	1 x	960 B	1 x
LB	INTEGER	VECTOR		STRUCT	

Accessed bytes:

- Offset 501760: Red (26), Green (550445), Blue (4500), Yellow (2 x 50445)
- Offset 41620: Red (41620), Green (10 x 50445), Blue (4500), Yellow (2 x 50445)
- Offset 41620: Red (41620), Green (3 x 50445), Blue (150445), Yellow (148445)

Contained XML data:

```
<File_write_at fid="0" time="1240651036.332362" count="501760" tid="1275069698" offset="501760" end="1240651036.332986" size="501760">
```

close

26445 550445

4500

2 x

2 x
550445

4500

960 550445

470 x

10 x

2 x

550445

41620

960 550445

...

1 Introduction

2 HDTrace

3 Visualization of datatypes

4 Summary

Conclusions

- HDTrace is a tracing and simulation environment
- Visualization of memory datatypes and I/O regions is important
- Sunshot enables to analyze datatypes in trace files
- Nested datatypes are visualized in a DAG
- Views are unrolled to indicate observed I/O access patterns

Future work

- Show collective (non-)contiguous I/O patterns in one window.
- Extract code into a small tool
 - Ease evaluation of datatypes (similar as MPIMAP)
 - Allow programming of datatypes (and I/O) in Java
 - Generate C code to create visualized datatypes