

PUBLIC

Sony Interactive Entertainment




**“Demo of a repository for statically compiled programs”
2016 US LLVM Developers’ Meeting**

Paul Bowen-Huggett
paul.huggett@sony.com

Agenda

Background




RFC

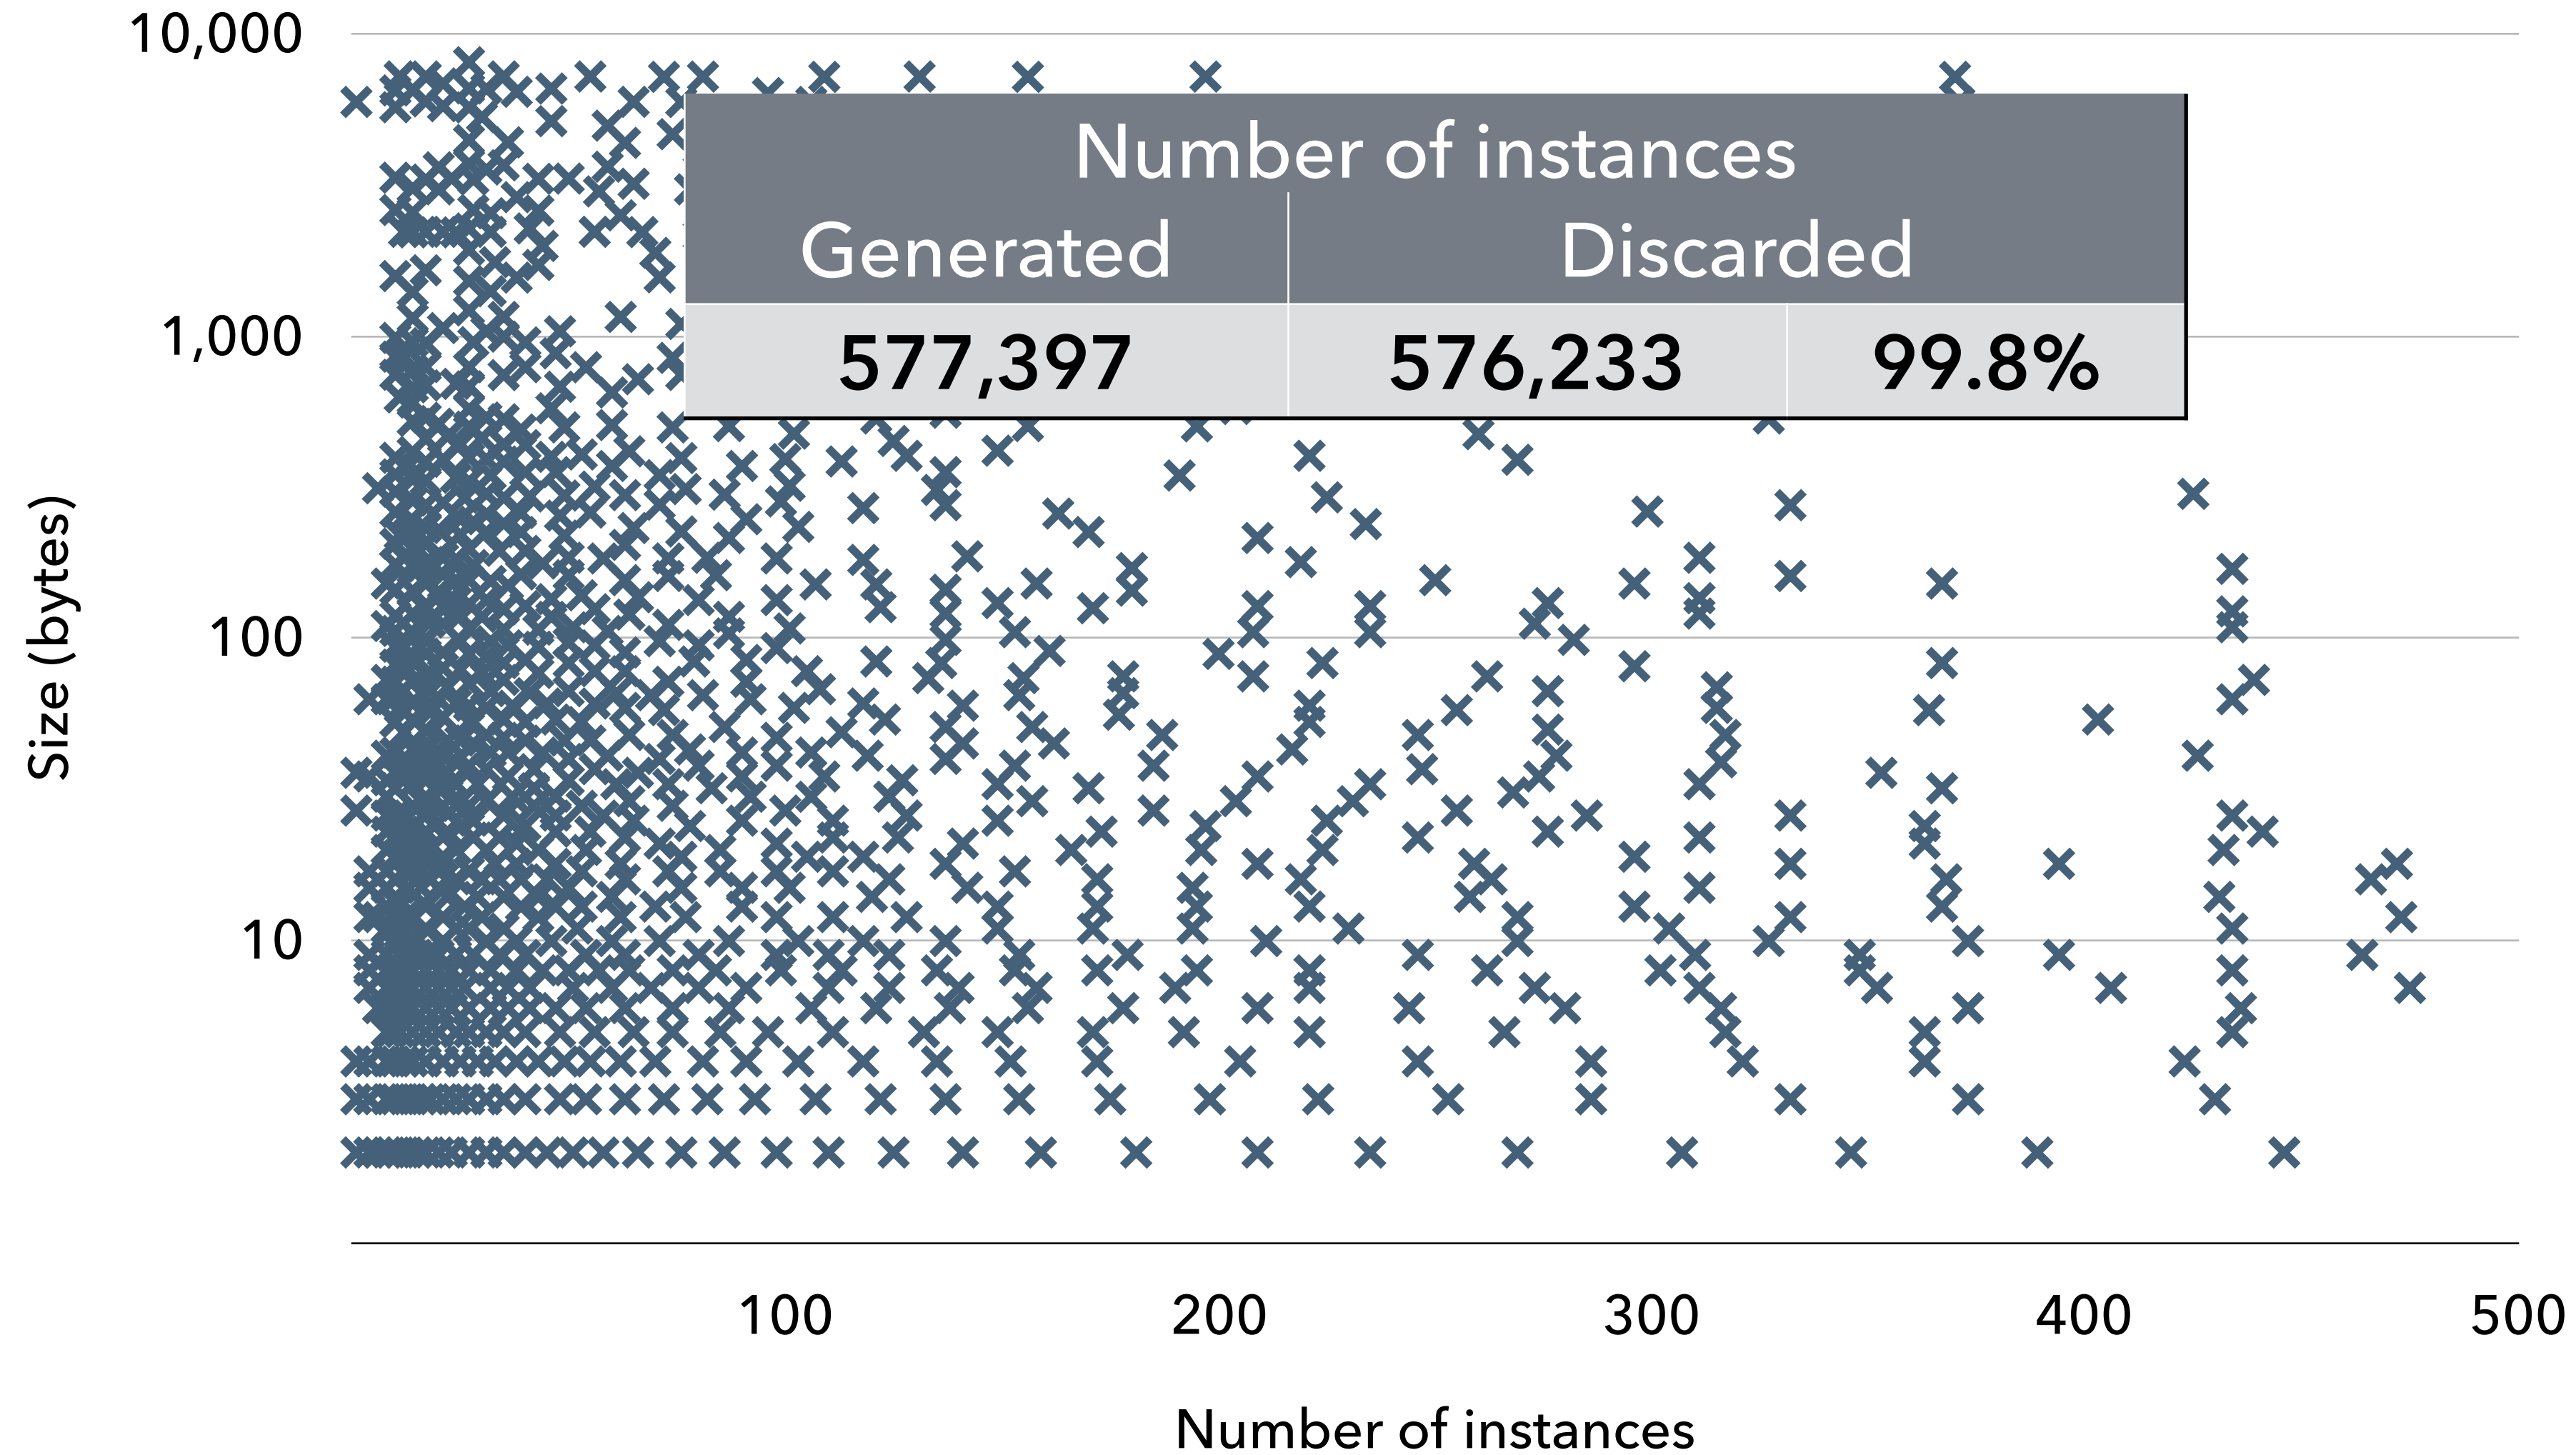
- Is the idea generally sound? Obvious improvements?
- Is it something we should think about for LLVM?
- There are several potentially related projects (C++ modules IFC, compilation database, ThinLTO, etc.) Views from respective owners?



Chromium Browser Build Ratios



Chromium Browser COMDAT Groups

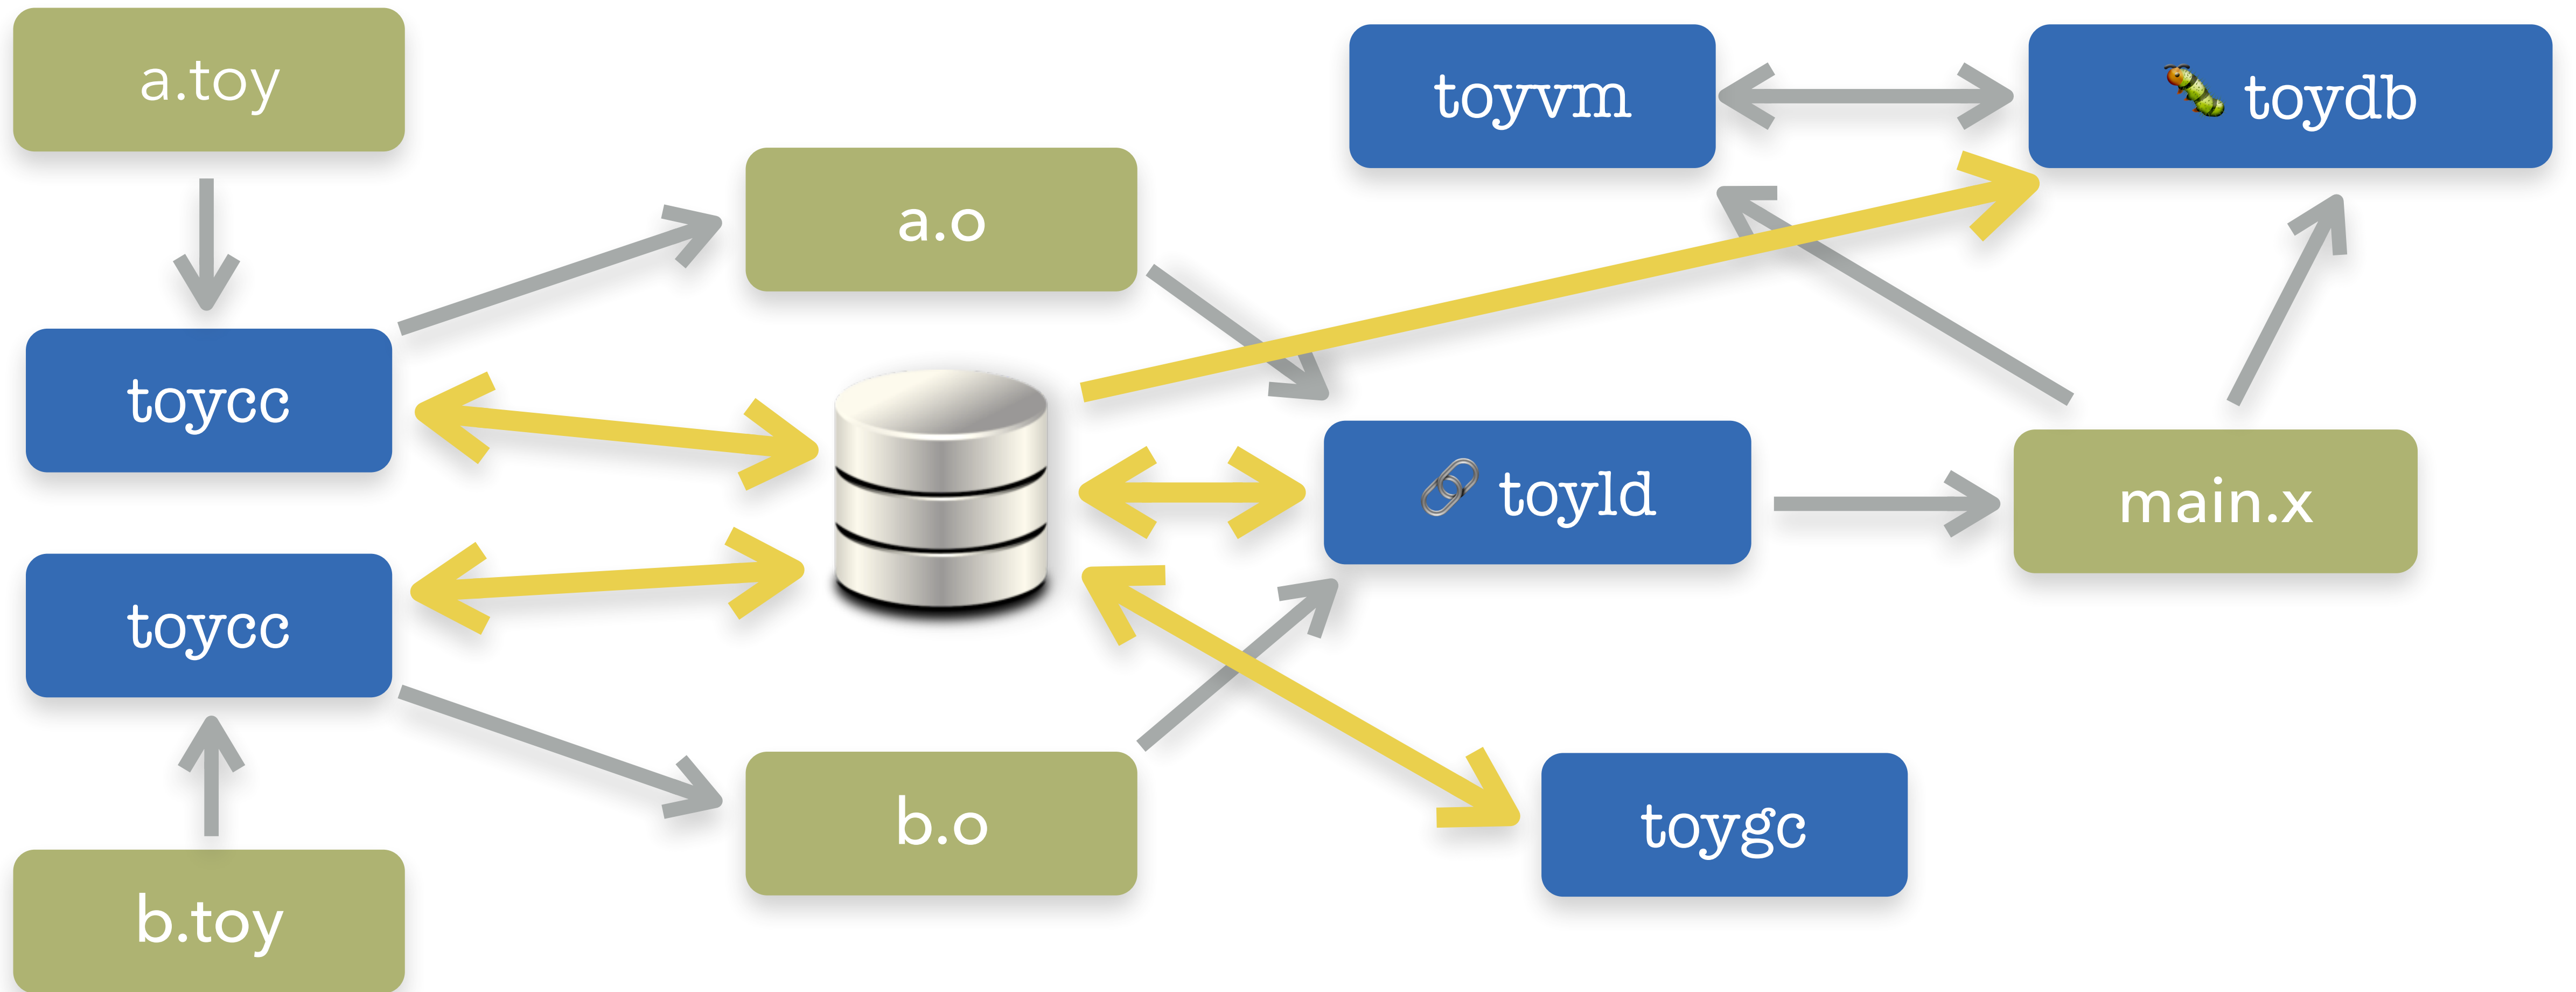


Toy Tools

- Toy programming language
- Available on github: <https://github.com/SNSystems/Toy-tools>
- Command line tools:

Role	Name
Compiler	toycc
Linker	toyld
Debugger	toydb
Runtime	toyvm

Role	Name
Garbage Collector	toygc
Strip	toystrip
Merge	toymerge



Limitations

1. It's just a toy!
2. Written in Python (3.5)
3. Output files are YAML
4. No concurrency
5. No backward compatibility
6. Says nothing about performance
7. The Toy language is nothing like C++:
 - VM has no registers, 3 stacks
 - Dynamic language, no user-defined types, no vague linkage...

Demo

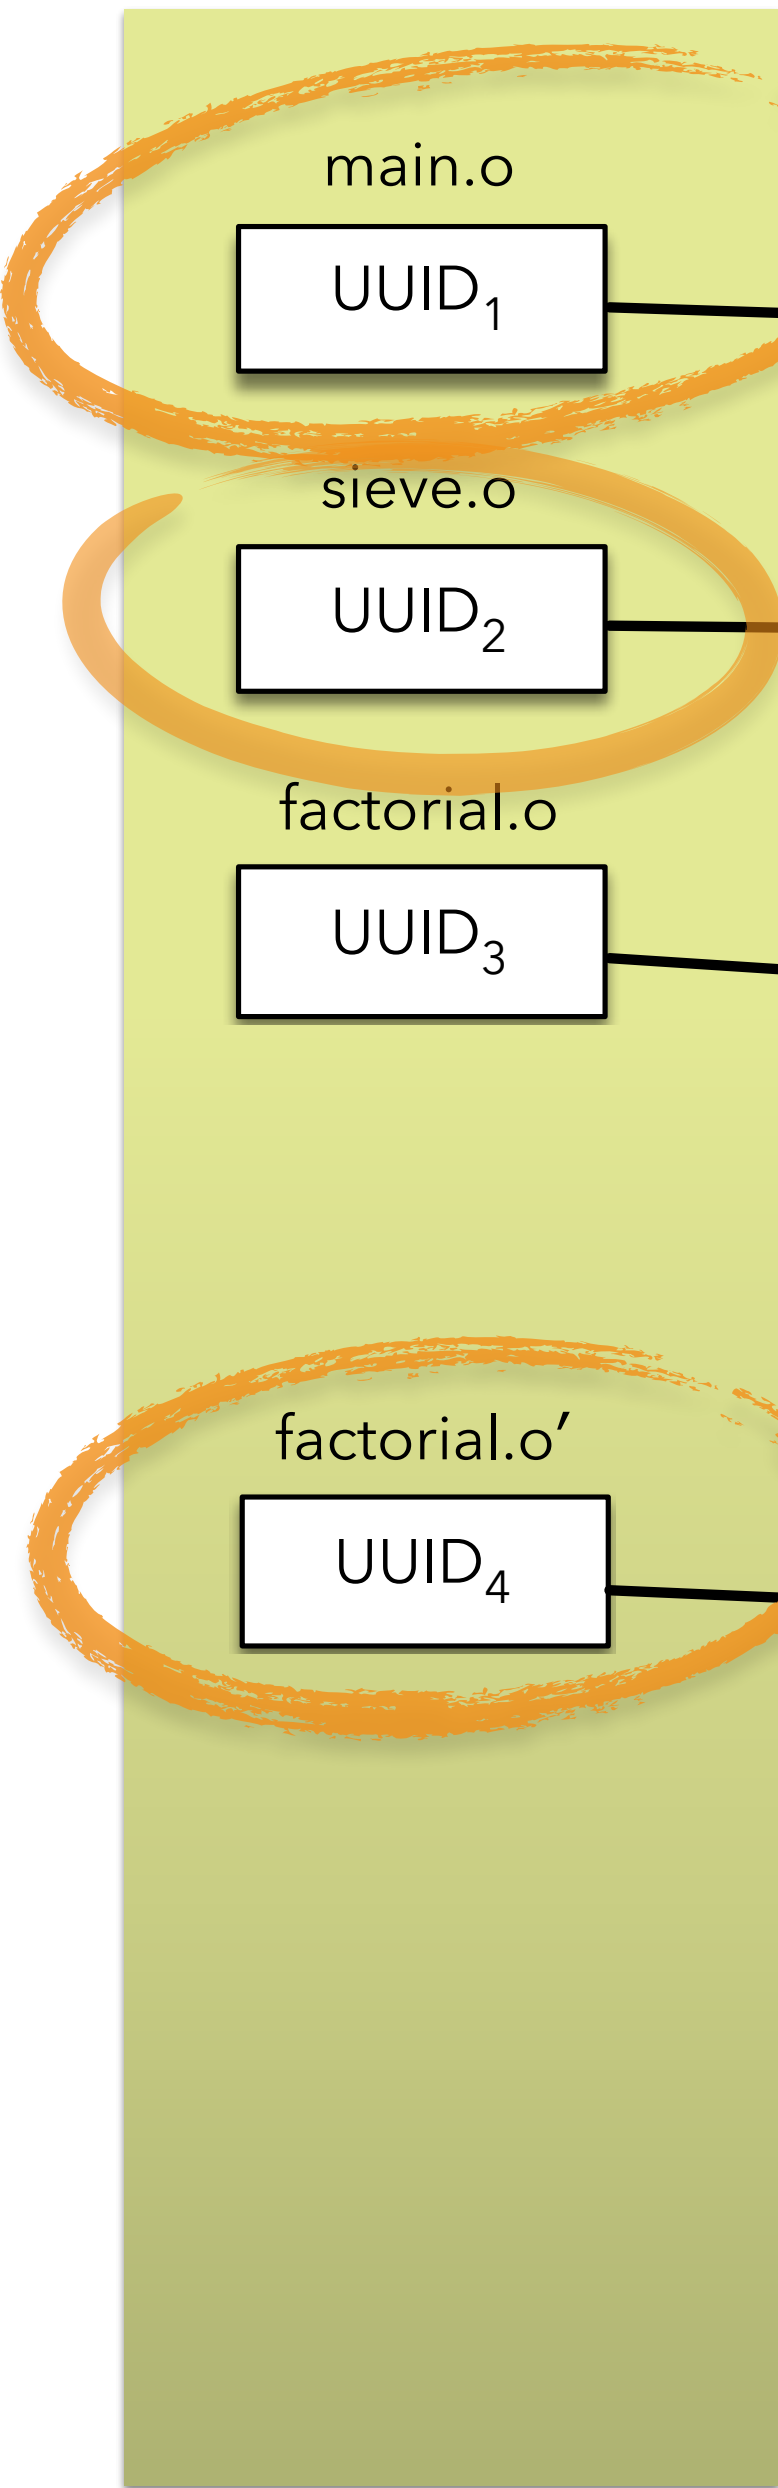
1. "Hello, World"
2. "Modules"
3. "Distributed"

"ticket" files

"tickets" table

entry point

"fragments" table



Key	Value
UUID ₁	name
	digest
UUID ₂	name
	digest
UUID ₃	name
	digest
	digest
UUID ₄	name
	digest
	digest

Key	Value
d(f ₁)	type
	binary
d(f ₂)	internal fixups
	external fixups
d(f ₃)	internal fixups
	external fixups
d(f ₄)	type
	binary
d(f ₅)	type
	binary

type	binary	internal fixups	external fixups
.text	55 48 89 e5 48 83	[]	"sieve"+0 "fact3"+0
.eh_frame	01 7a 52 00 01 78 10 01	.text+0x19 .text+0-x2f	[]

type	binary	internal fixups	external fixups
.text	55 48 89 e5 48 83	[]	"factorial"+0

type	binary	internal fixups	external fixups
.text	66 4e 89 e5 48 83	[]	"factorial"+0

name

digest

name

digest

name

digest

name

digest

name

digest

type

binary

internal fixups

external fixups

type

binary

internal fixups

external fixups

type

binary

internal fixups

external fixups

type

binary

internal fixups

external fixups

type

binary


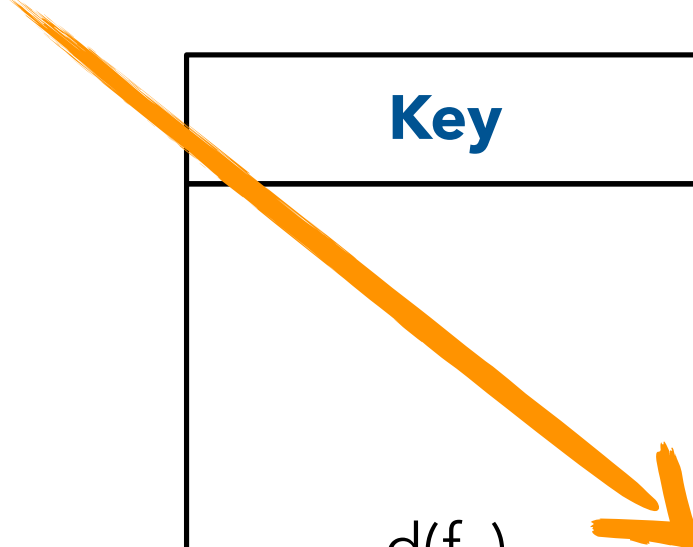
internal fixups

external fixups

entry point

"fragments" table

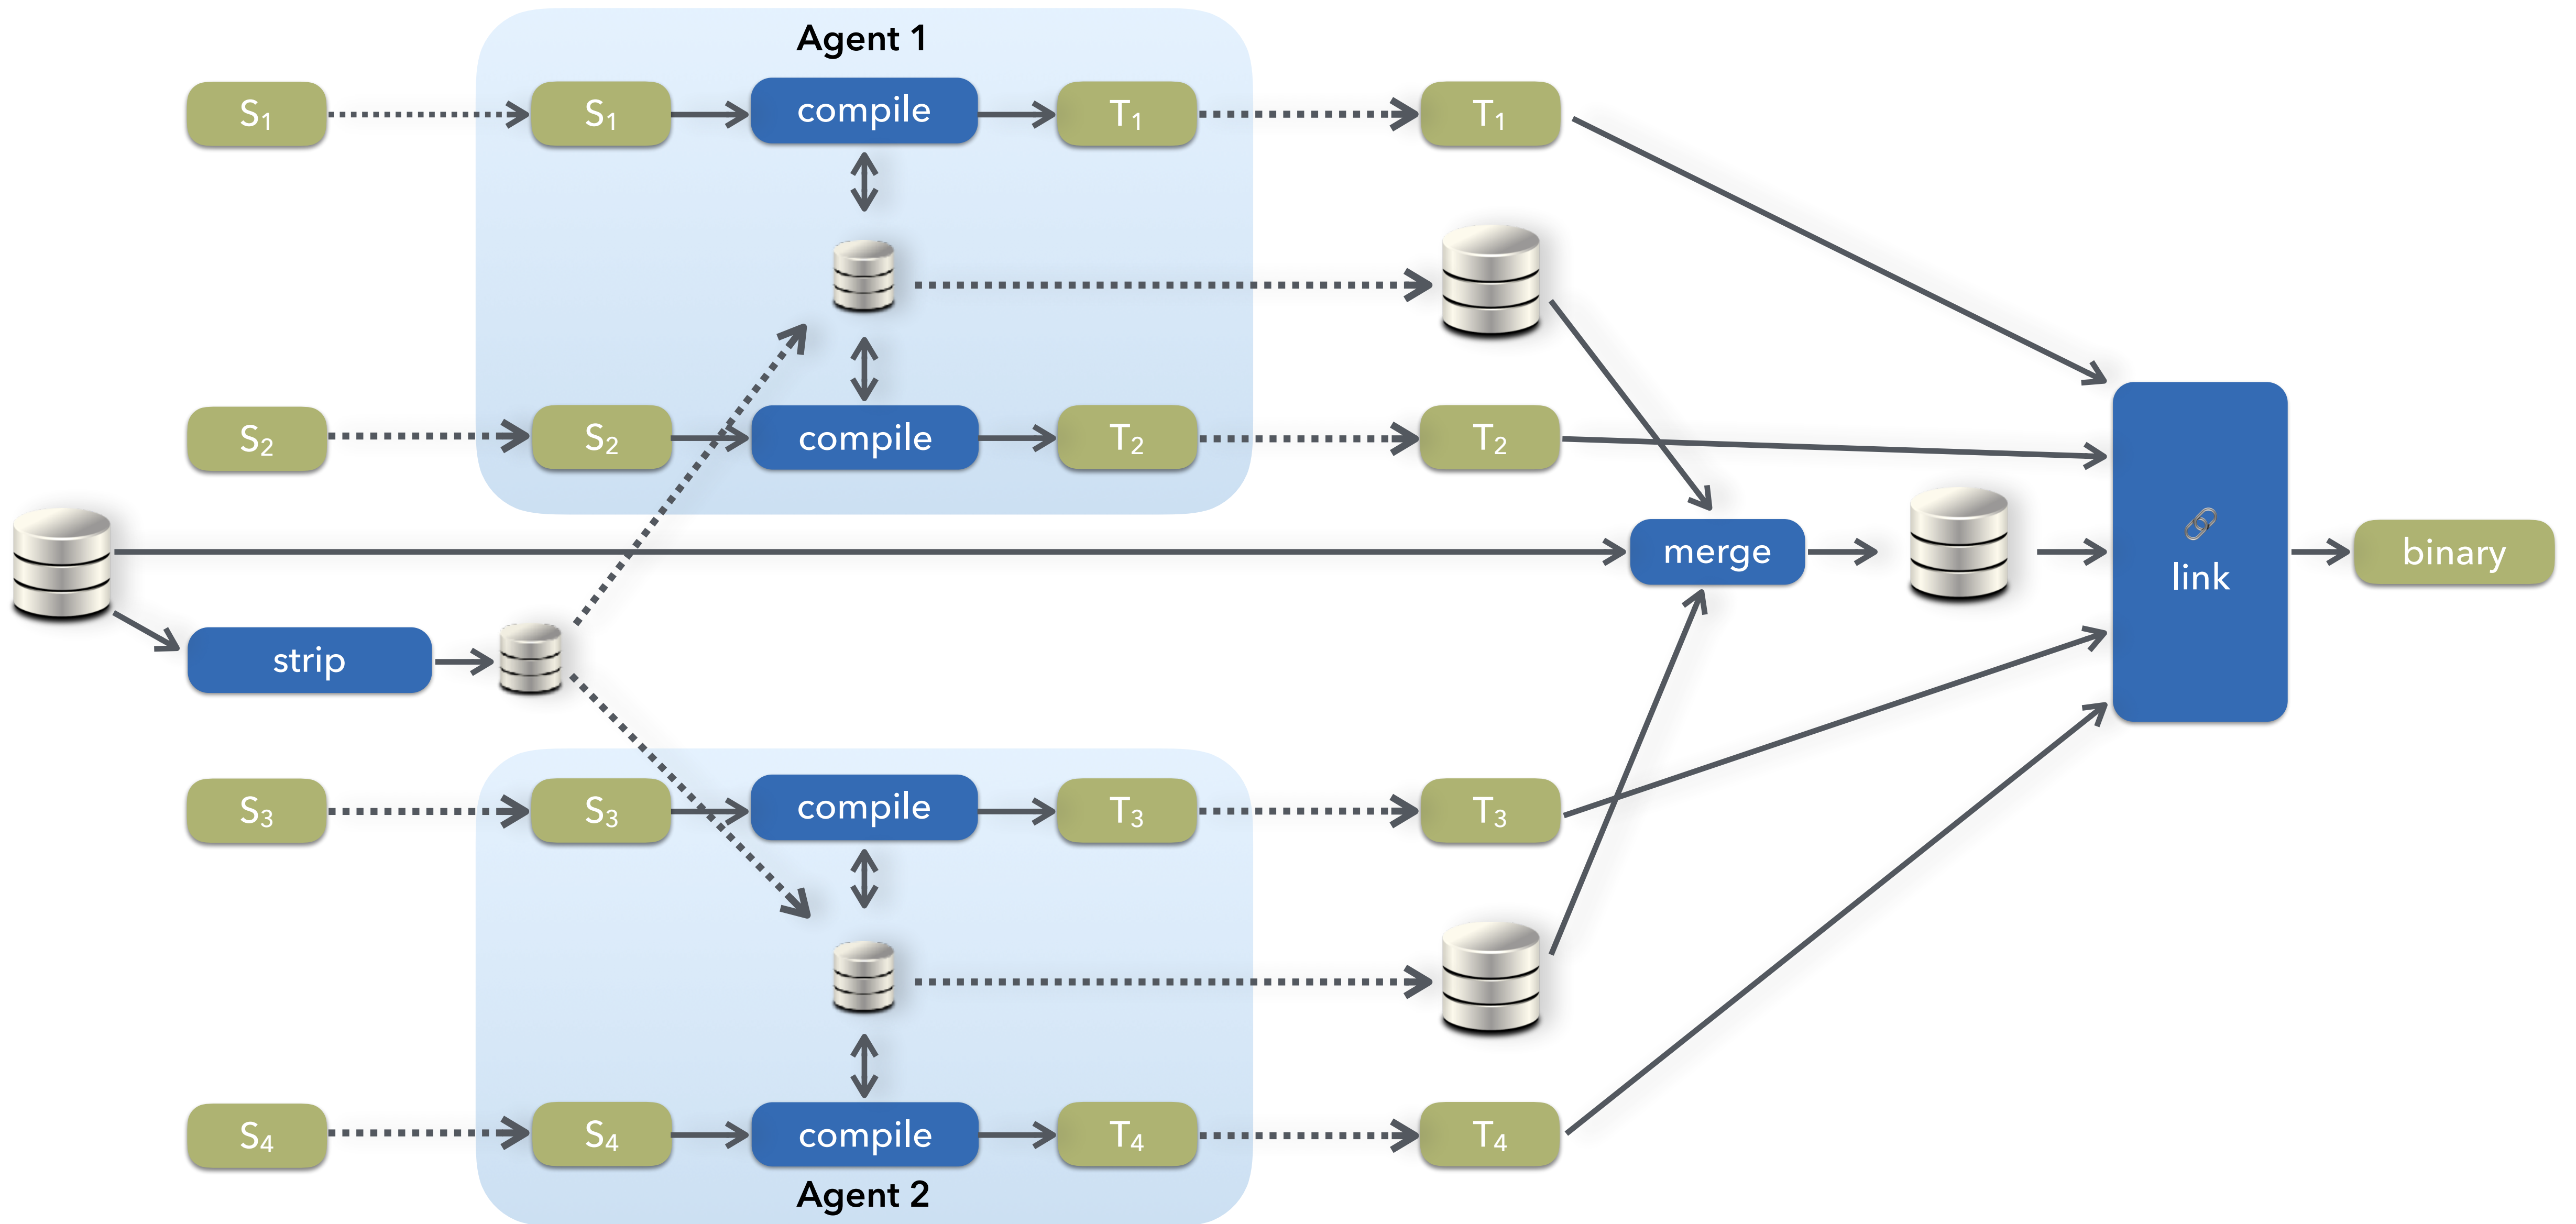
Key	Value			
d(f ₁)	type	binary	internal fixups	external fixups
	.text	55 48 89 e5 48 83	[]	"sieve"+0 "fact3"+0
d(f ₁)	.eh_frame	01 7a 52 00 01 78 10 01	.text+0x19 .text+0-x2f	[]
	d(f ₂)	...		
d(f ₃)	...			
d(f ₄)	type	binary	internal fixups	external fixups
	.text	55 48 89 e5 48 83	[]	"factorial"+0
d(f ₅)	type	binary	internal fixups	external fixups
	.text	66 4e 89 e5 48 83	[]	"factorial"+0



Distributed Builds

Should the repository be a network resource?

Distributed Build



Challenges?

- Remember, it's just a toy... Need a production-ready C++ implementation
- Understand real-world growth rates and GC strategy
- Doesn't show solutions to:
 - Fast storage with efficient indices
 - LLVM IR hashing
 - Efficient debug type references

Conclusion

- Potential to reduce re-compile times by ~60% ("speed-of-light" based on Chrome Debug)
- Small code changes benefit the most
- No source code changes
- Eradicate duplication and redundancy **at source**: minimize link-time processing and copying
- (Almost) No change to workflows
- Next steps:
 - Data store (in-process, memory-mapped hash tables)
 - Prototype:
 - IR hashing
 - MC back-end
 - Repository-based linker

Q & A