

lld

Solving the Linking Performance Problem

lld is a system linker for multiple platforms

Windows

Linux

Darwin

FreeBSD

...

Ild is a cross linker. Always.

lld is free software under the UoI/NCSA License and part of the LLVM project

Status

Self hosts

x86-64 Linux

libstdc++

libc++

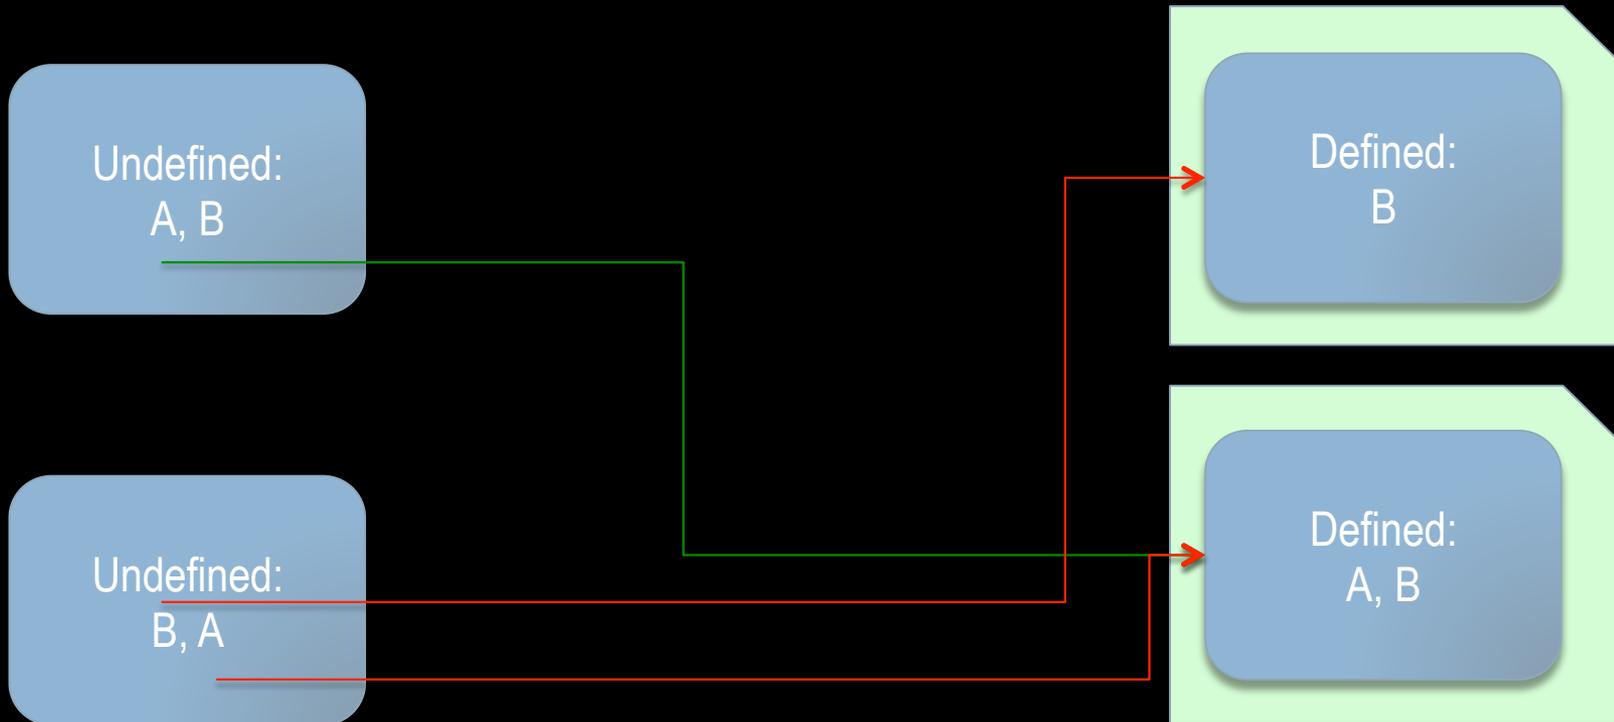
Status

Performance on par with gold on Windows

Why is linking slow?

Complex semantics

The AB-BA problem



Complex semantics

Command line order

--{start,end}-group

-z rescan

-z rescan-now

...

Input formats not optimized for linking

ELF is optimized for runtime linking and loading

Archives

How does Ild make it fast?

Speculatively Resolve

Read in parallel

Not semantically correct

Mostly correct

Fast

Sequentially Resolve

Resolve again

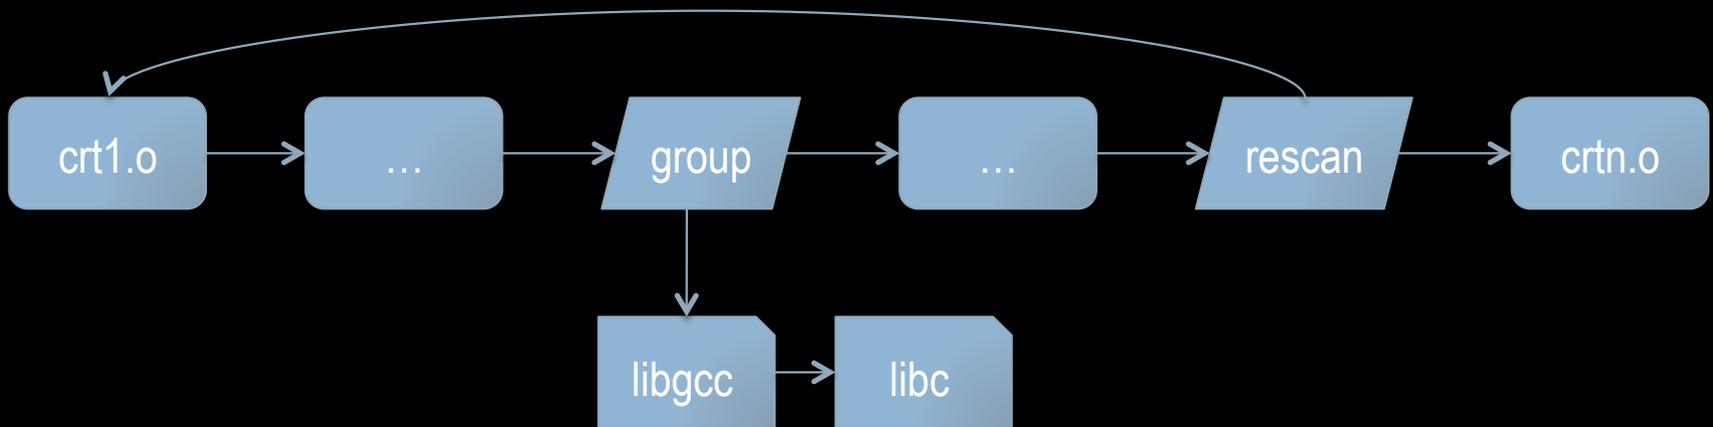
Semantically correct

Use cached object loads

Load what we missed

Most lookups hit the cache

Input file graph



Faster?

Change How We Link

Simplify Semantics

Command line order + archive extraction force sequential resolving

Ideal is resolve in parallel

Achieved by linking everything on the command line

Pick resolved symbol based only on command line order and weakness

New Format

Reduce the number of memory pages we have to touch

Data needed to resolve is highly compacted

Questions?

Michael Spencer
bigcheese@gs.com