



DESIGN

FAST
SERVER

BLOOP

CONCURRENT

BUILD

CHALLENGES

VEGAN

ORGANIC



ENERGY FRIENDLY

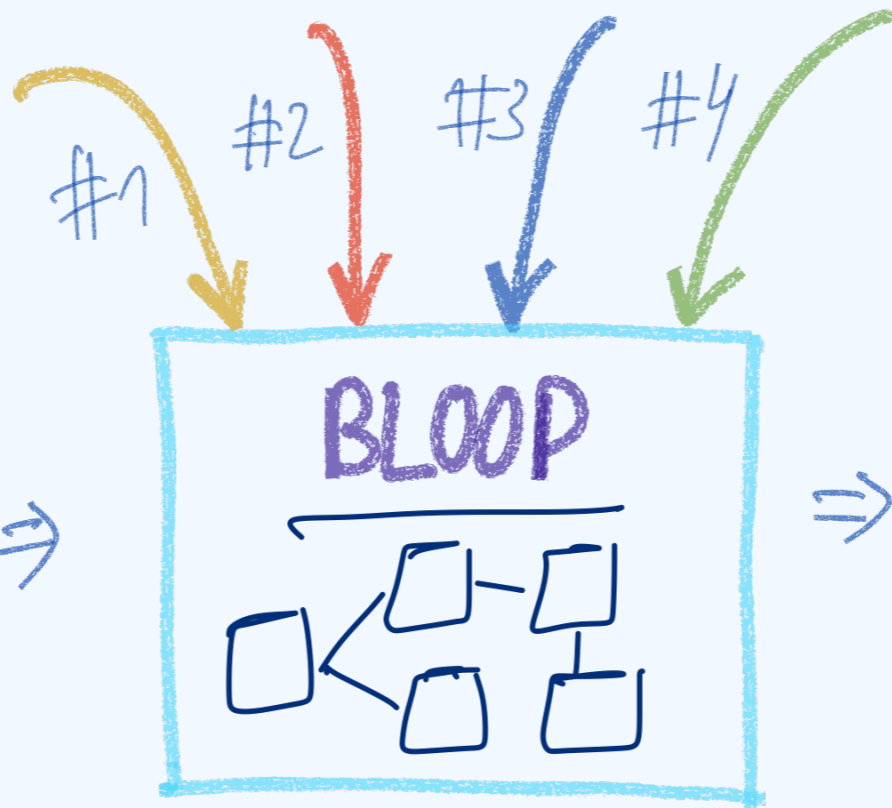
SUCH FAST

I RUN IN THE BACKGROUND OF YOUR MACHINE SERVICING REQUESTS AS FAST AS POSSIBLE. I WANT TO MAKE YOU **HAPPY**.

REAL FAST

`bloop { compile, test, run, debug }`


`$WORKSPACE/.bloop/*.JSON =>`



\Rightarrow

DATA
+
SIDE EFFECTS

Highlights



Tooling as a service, focus on protocols < BSP
Nailgun

Integration-friendly < high quality docs
Tooling → launcher < binary library

Integrates with editors, build tools, third-party tools.

A nice trade-off between centralization & decentralization,
in an environment with many constraints & particular culture.

Credits



Scala Center (Heather Miller, Sébastien Doeraene)

Martin Dohem, co-author of Bloop

Olafur Páll Geirsson, creator of Metals

Tomasz Godzik, Marek Żarnowski @ Virtuslab

Developers solving similar problems (Eclipse, IntelliJ, ...)

External contributors & happy users

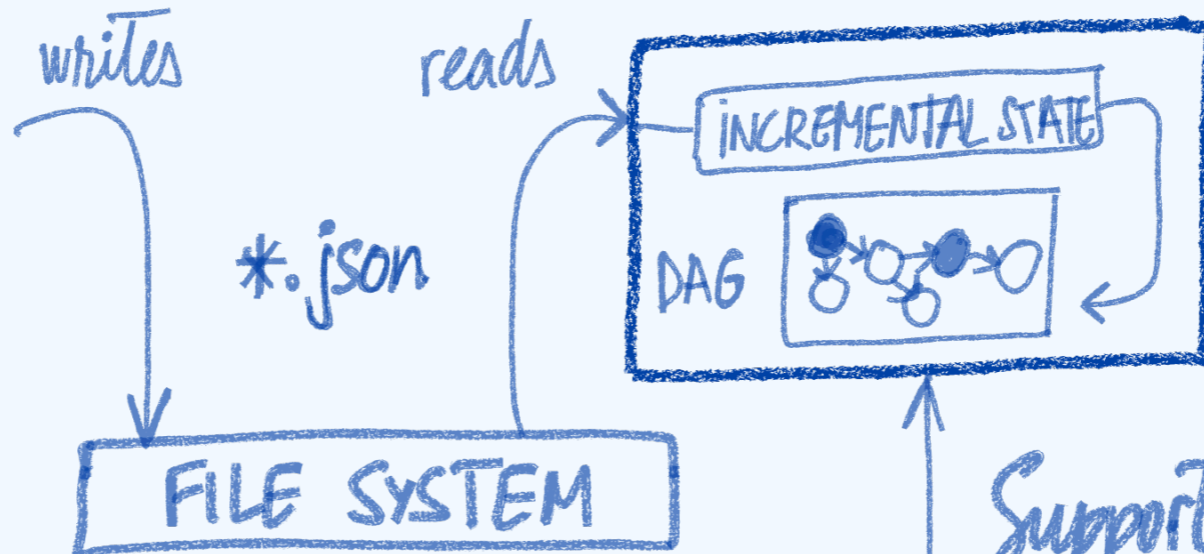
How does bloop understand your build?

Build tools



BLOOP

daemon



FILE SYSTEM

Supports:

- Fast import project
- Updates on build changes

bidirectional

1 USER



HIGH EXPECTATIONS #1
WANTS TO BE BLOCKED BY ITS THINKING #2
LAZY, DOESN'T READ DOCS #3

independent clients

N CLIENTS = $k \cdot m$

k WORKSPACES

m TOOL INSTANCES



m

WORKSPACE #1

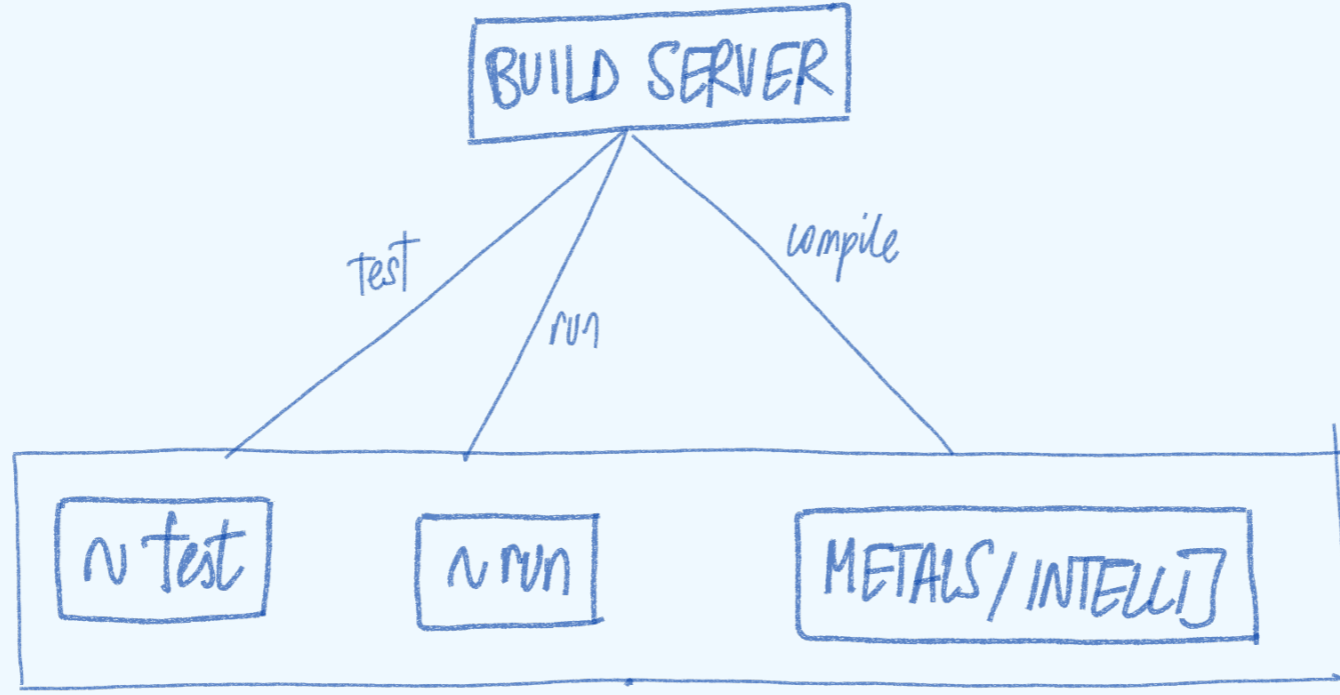
WORKSPACE #N

1 BLOOP



Case study: concurrent clients.

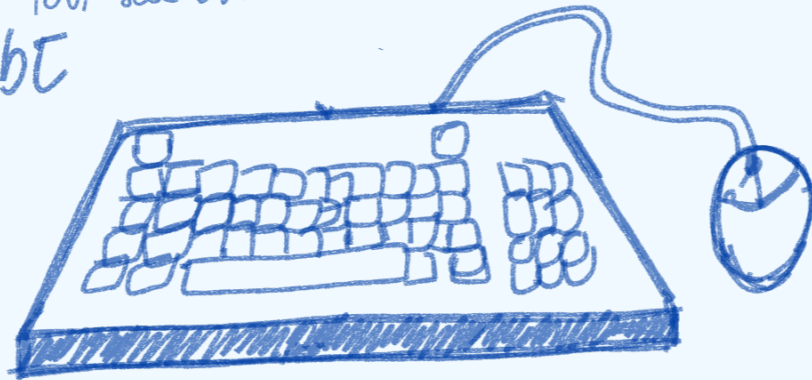
currently not possible!




at the same time

concurrent clients.

2 build tool sessions
sbt



useful to work on web applications!

Semantics of a build server?

What if...

- Two clients compile at the same time?
- clients cancel compilation?
- A client compiles when another is running tests?

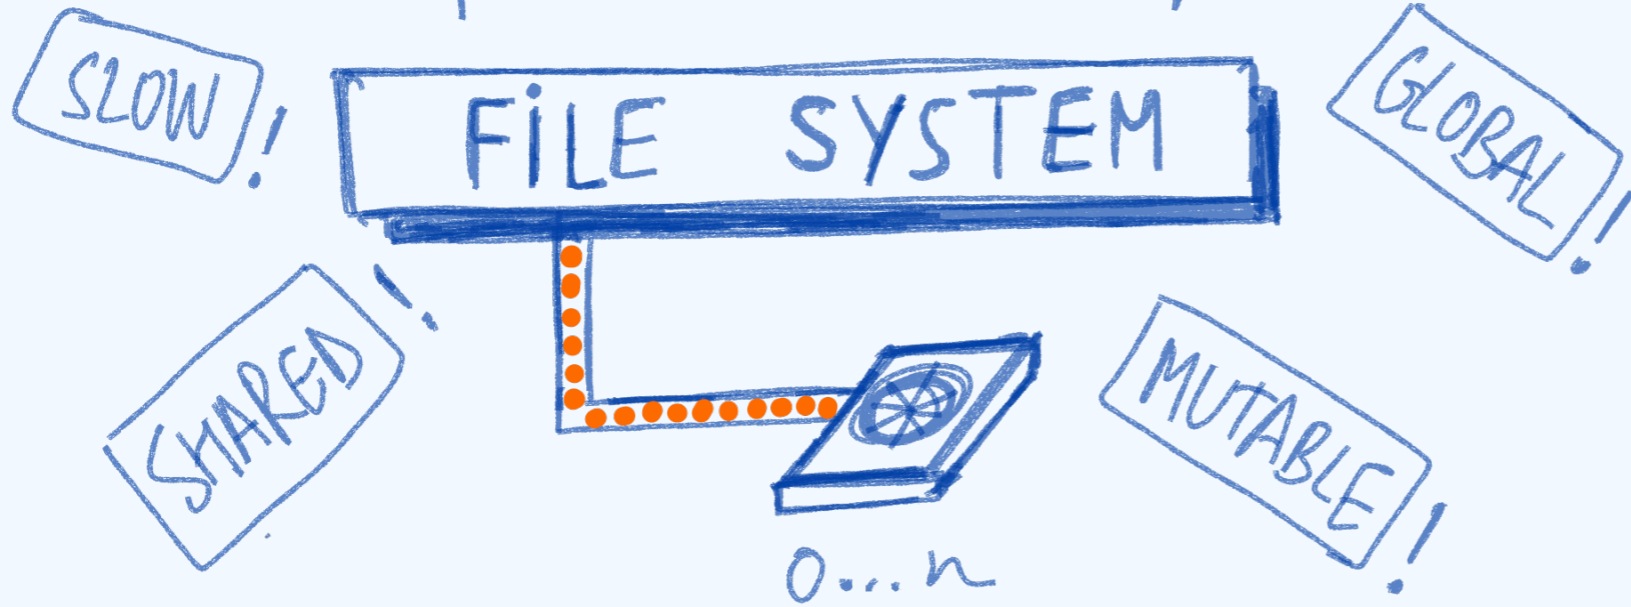


No previous state-of-the-art, ñ
oh no!

Compile Inputs (low-level)

- sources (*.scala, *.java)
- resources (generated or non-generated)
- classpath (provided by build tool)
- scala version and options
- project metadata (more generally)
- internal state.

Most inputs live in the file system...



oh no!



2 properties to rule them all


Compile deduplication

same inputs = same outputs
requires emulating a compilation
motivated by efficiency



Compiler isolation

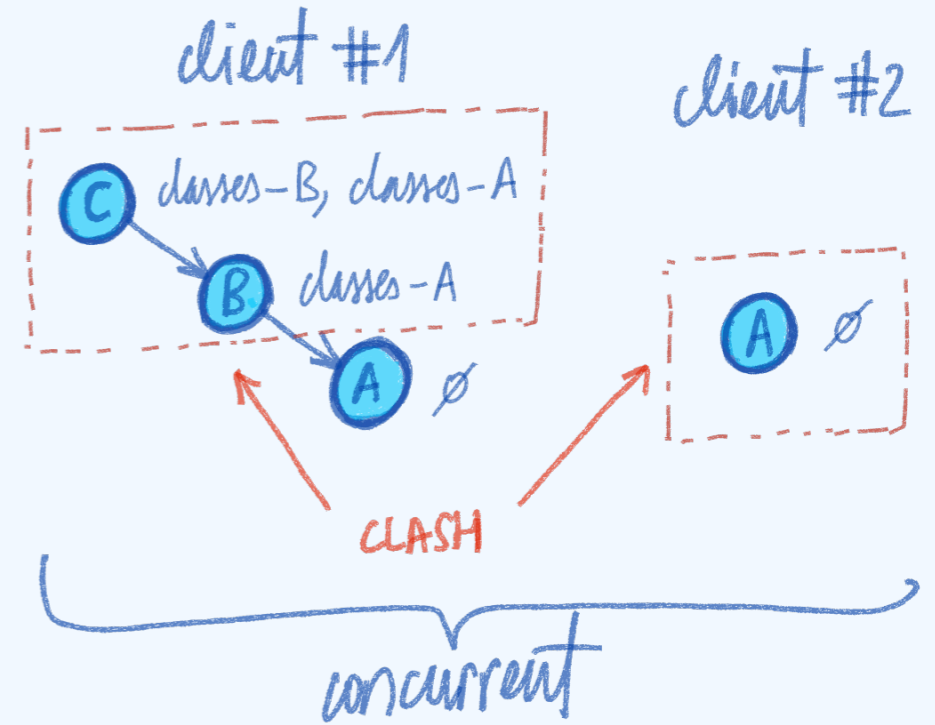
different inputs \Rightarrow different outputs
requires inputs to be $\left\{ \begin{array}{l} \text{immutable if shared} \\ \text{mutable if unique} \end{array} \right.$
motivated by correctness



Compiler isolation

Challenges

- Compiling overlapping DAG subgraphs
Classpath contains references to current dirs.



- Accessing latest compiler state in server



RACE CONDITIONS!

- Persisting latest state of a project in-disk.

Sounds familiar?


We need

- Atomicity
- Consistency
- Isolation
- Durability

Solutions for compiler isolation



Blocking { waste of resources, doesn't scale
nobody likes unnecessary blocking ☹



Non-blocking

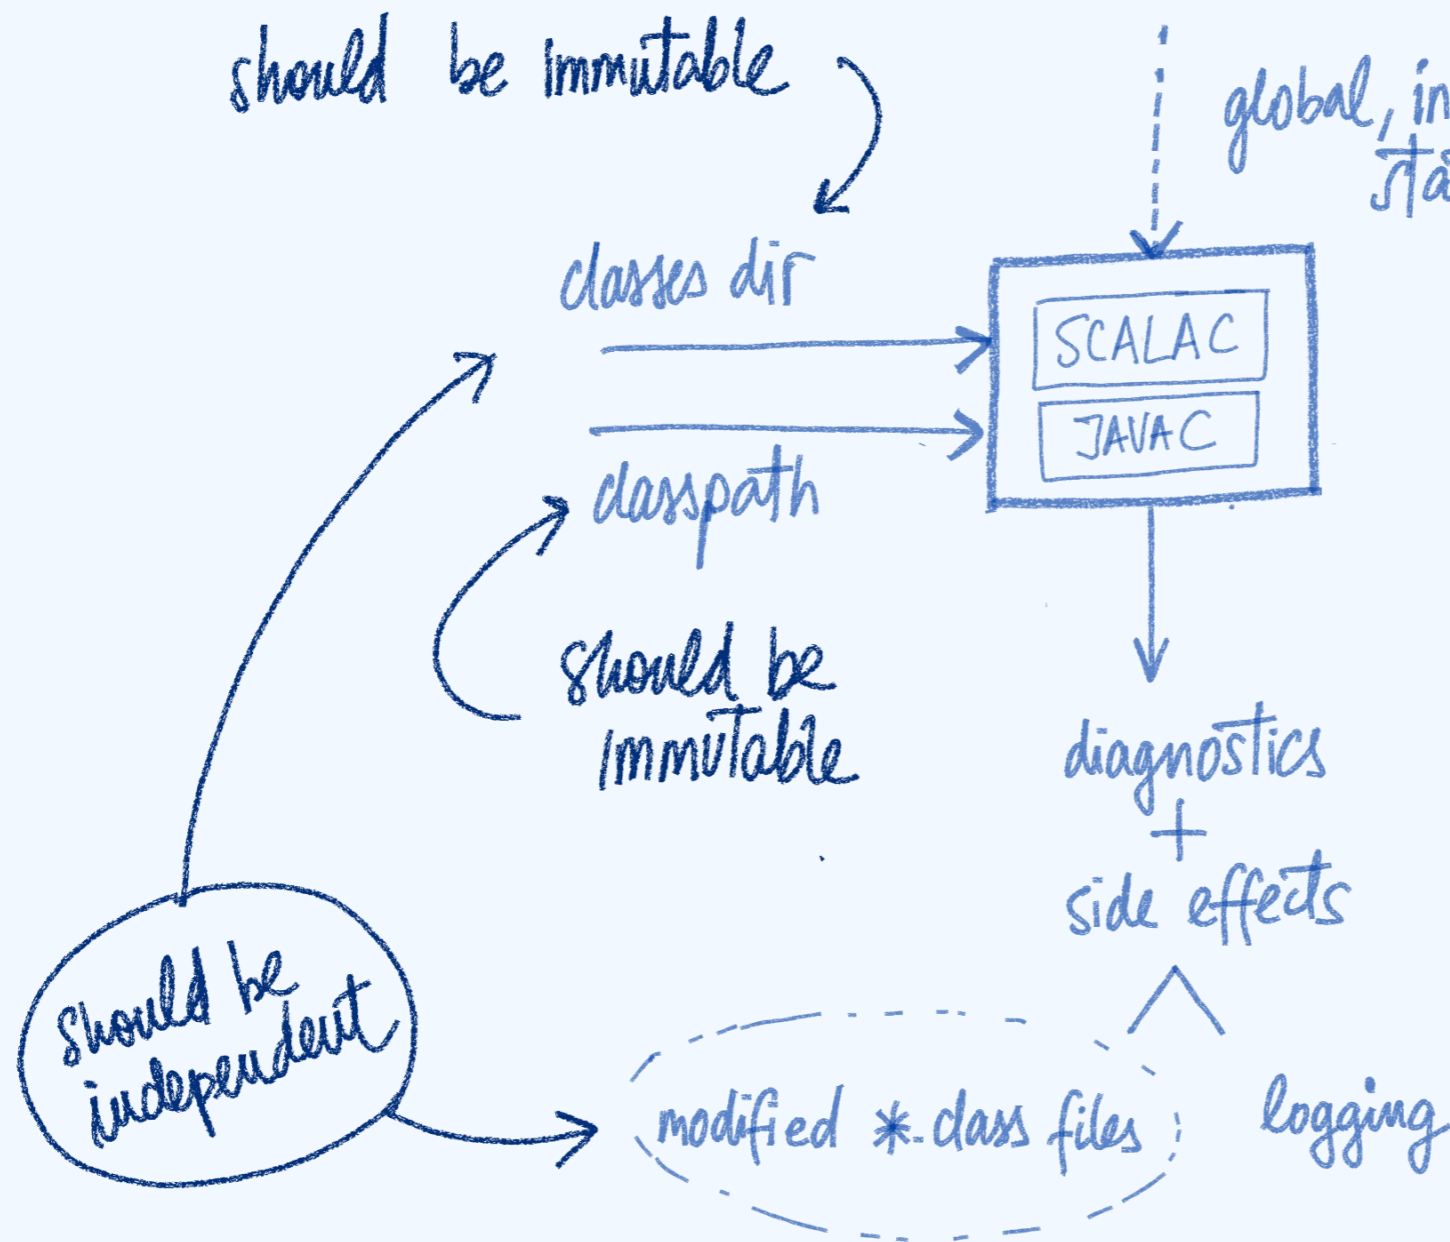
difficult but worth it!

} devise clever way to avoid sharing
use transaction-based compilation engine to provide ACID semantics
implies independent outputs
⇒ make approach fast

Avoid sharing state < files → compilation outputs
in-memory mutable data



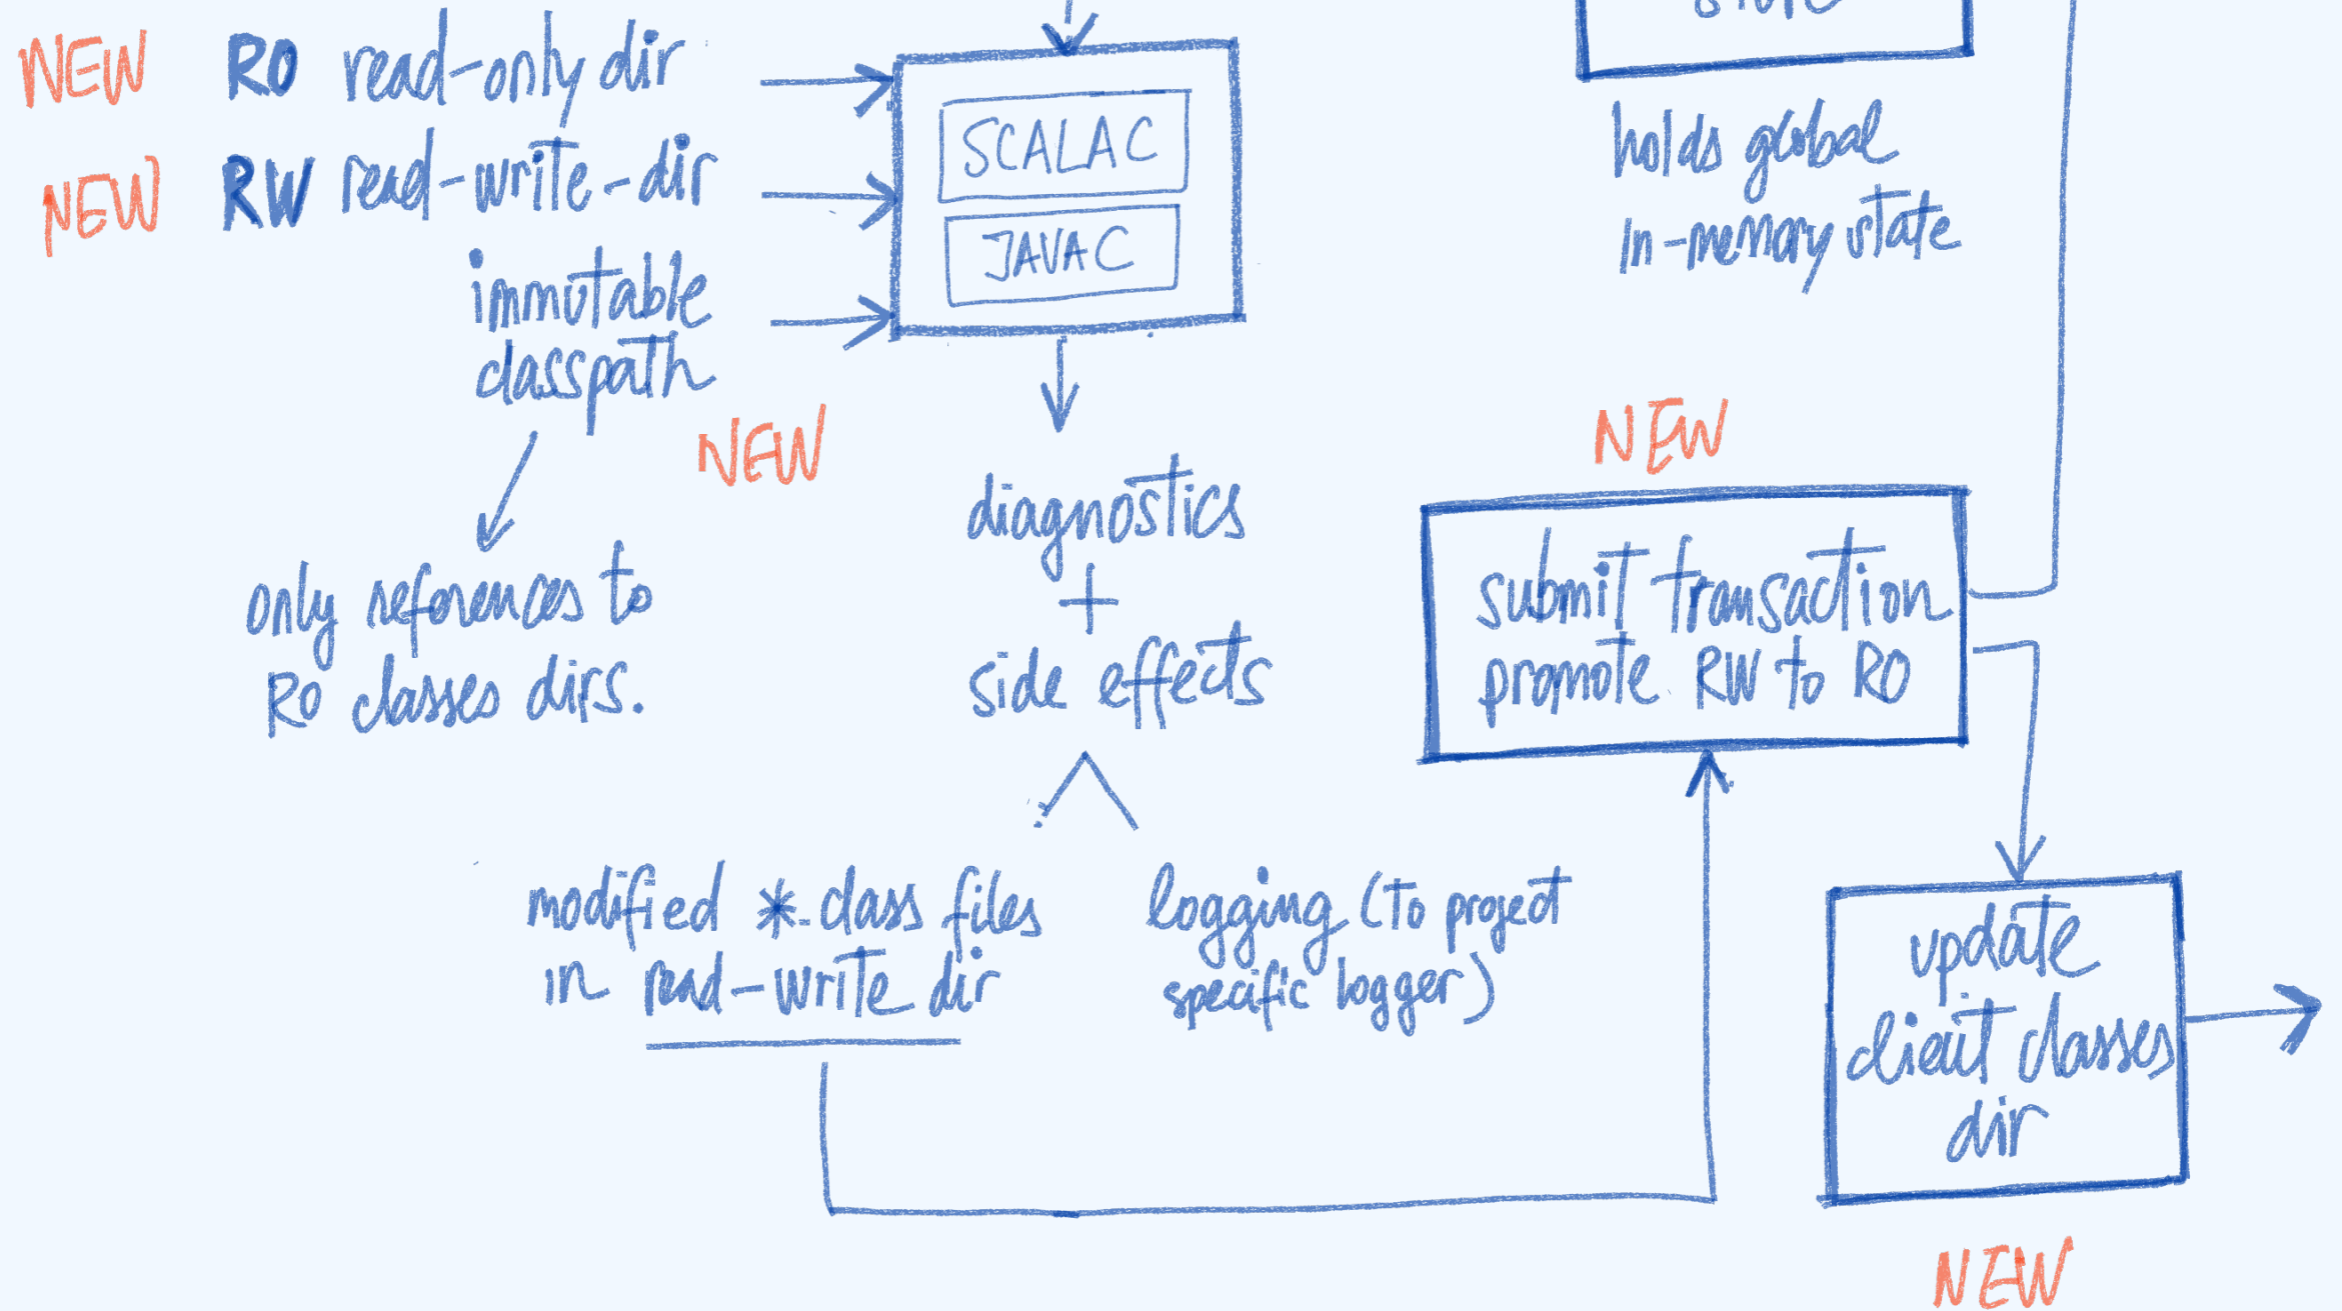
Avoid sharing state < files → compilation outputs
in-memory mutable data



read & writes should be controlled to avoid races

our handling of directories must change!

New design



Compiler isolation

Every client has its own unique set of directories.
They get populated by the compilation engine (in parallel).

Build tool generating build files keeps the original dirs.

Compilation engine works with read-only directories
New read-write directory promoted to read-only when result is cached.

Compile deduplication

Avoids extra compilations when inputs are the same.

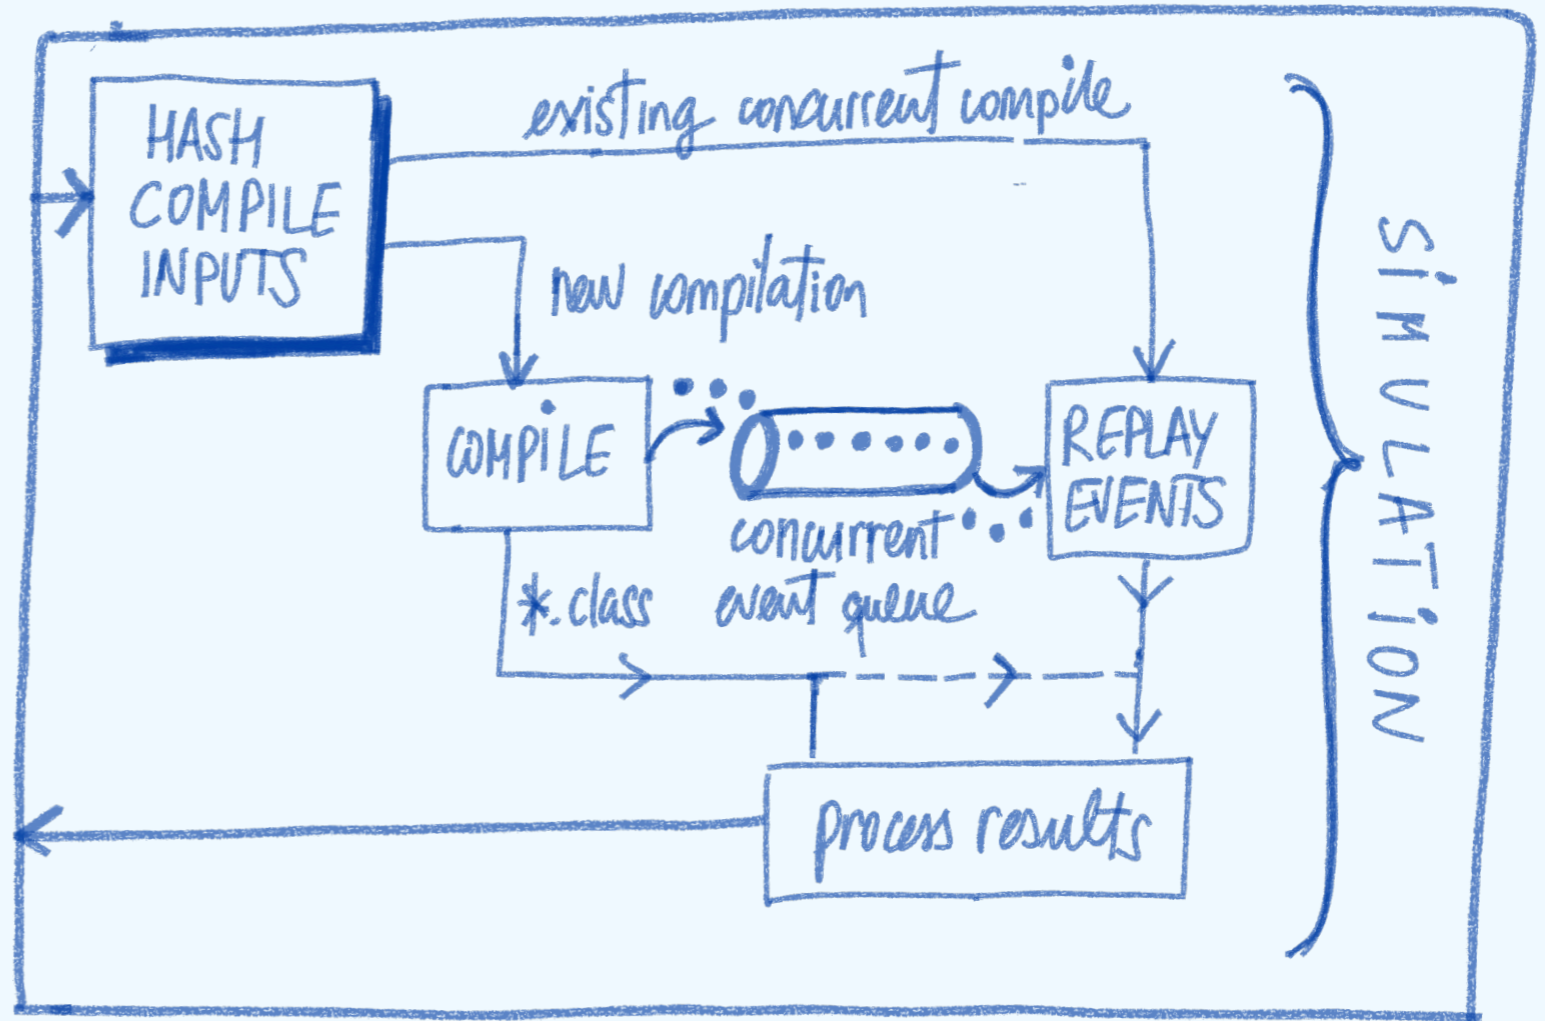
— Crucial for real-world use cases: several clients sending many compile requests for the same build at the same time

Key idea = simulate compilation

Compile deduplication

All clients should perceive the same behavior

Simulation depends on client data



Offloading compilation from sbt

sbt shell

> test

> ~ run

(sbt) compile

make sure bloop configs
are up-to-date

ask bloop to compile via BSP

Seamless integration means...

- people don't need to change their workflows
- compilations are shared with other clients.



JUST WORKS
TECHNOLOGY


Testing the build server

Treats build server as a blackbox and tests all build side effects.

Focus on compiler invariants before & after build actions.

- Compiler diagnostics, compilation products (+ metadata), etc.

Randomized tests of concurrent clients simulating real-world scenarios
~ 11 KLOC of tests for (protocol, platform, client type).




Performance evaluation

Basic numbers.
(for macOS)

~ 200 - 500 ms incremental compiles
in medium-sized codebases

~ 500 ms full no-op compiles
in big codebases + 200 modules



Concrete numbers

AKKA

sbt	2 minute, 25 seconds
bloop	1 minute, 2 seconds

CIRCE

sbt	1 minute, 5 seconds
bloop	50 seconds

Oracle Java 8
linux, Intel i7 @ 3.5GHz

GUARDIAN / FRONTEND

sbt	44 seconds
bloop	29 seconds

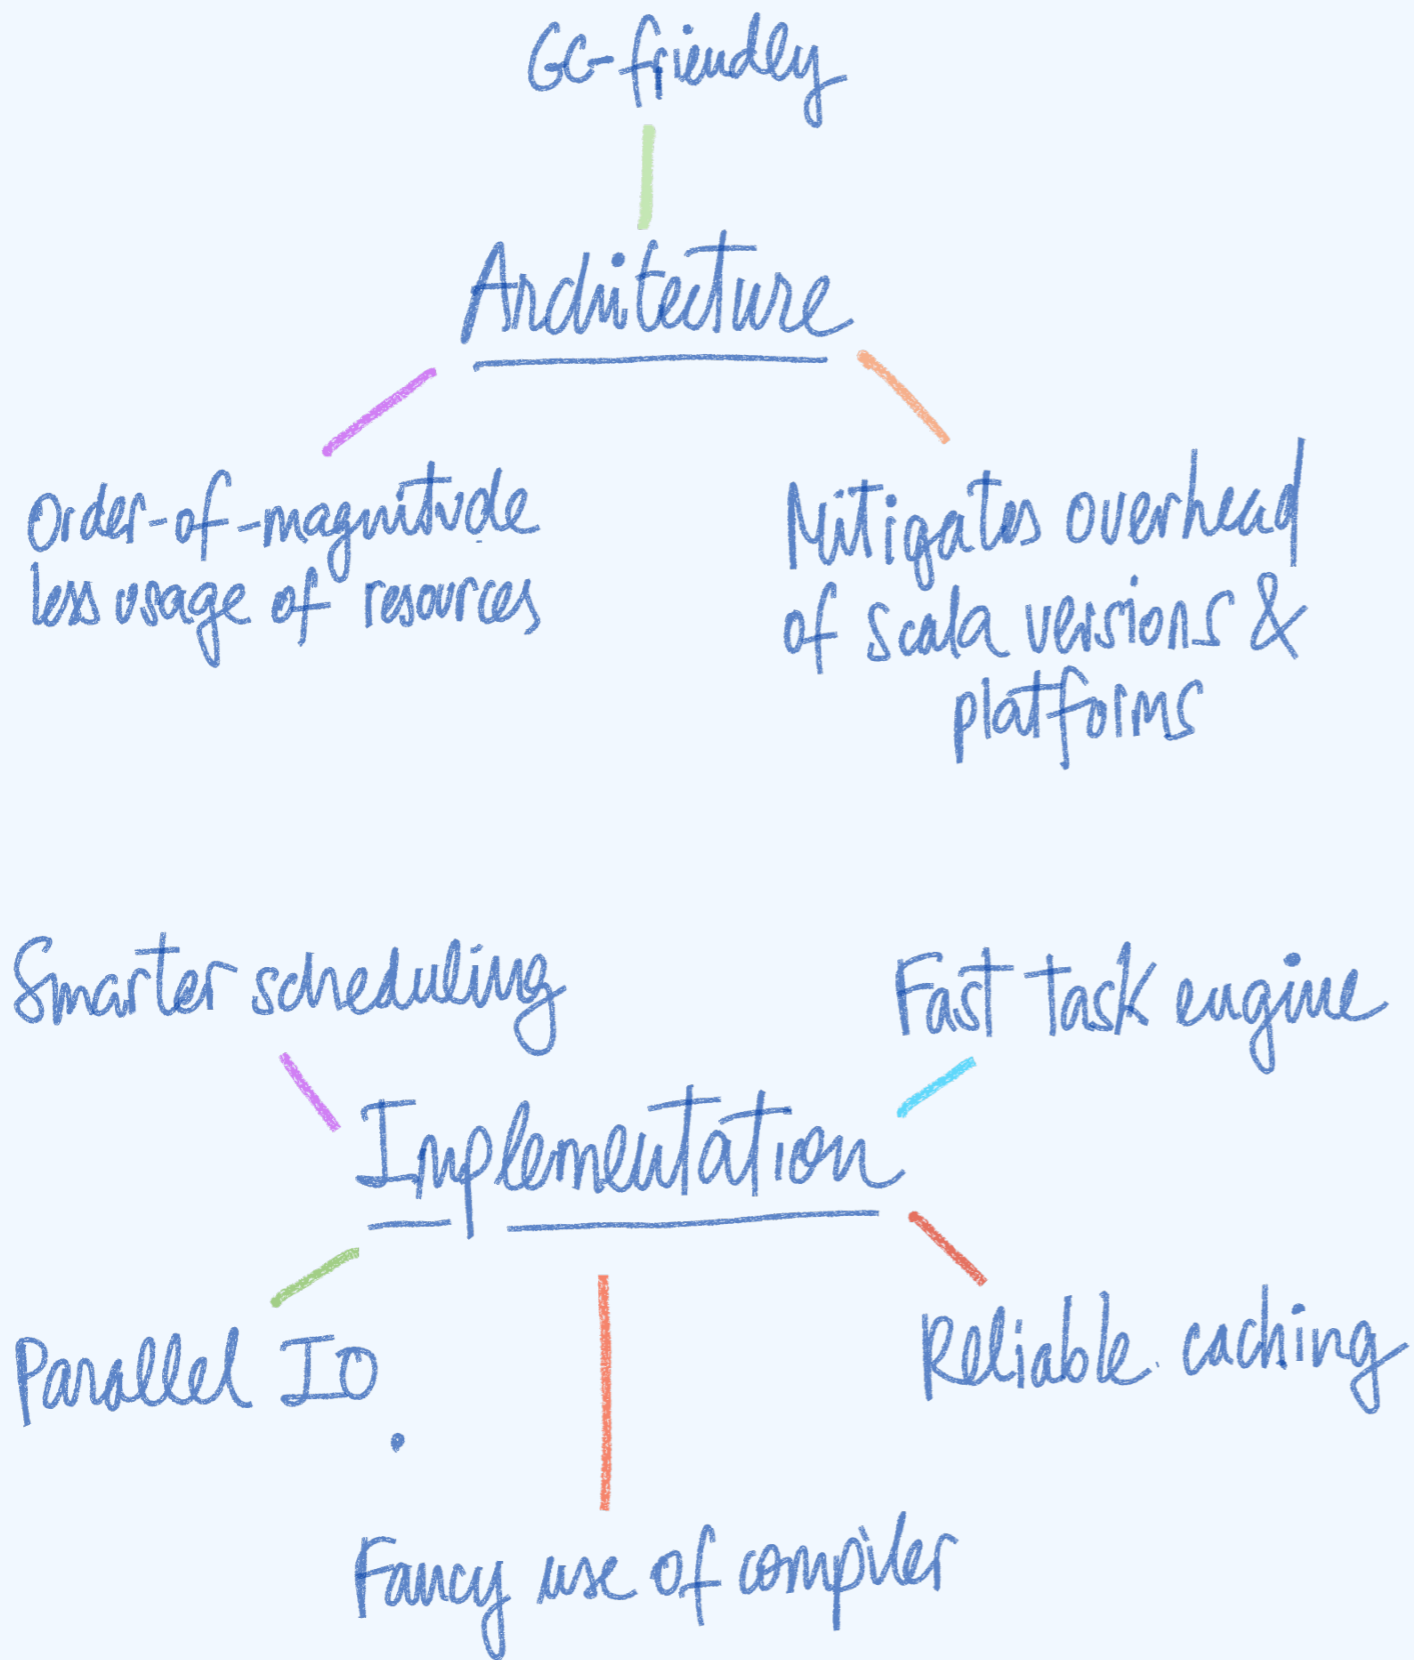
HTIP4S

sbt	48 seconds
bloop	43 seconds

ATLAS

sbt	9 seconds 90ms
bloop	9 seconds 950ms

)!




Where does Bloop speed come from?
(even if it does more work?)

Caching

Caches compile inputs heavily \leftarrow Sources
Classpath
But reliably: no dependency on file watching

Caches compiler plugins classloaders automatically
upcoming: will also cache macro classloaders

Bloop's good defaults policy



"one server to rule them all"
vision

users



win
win

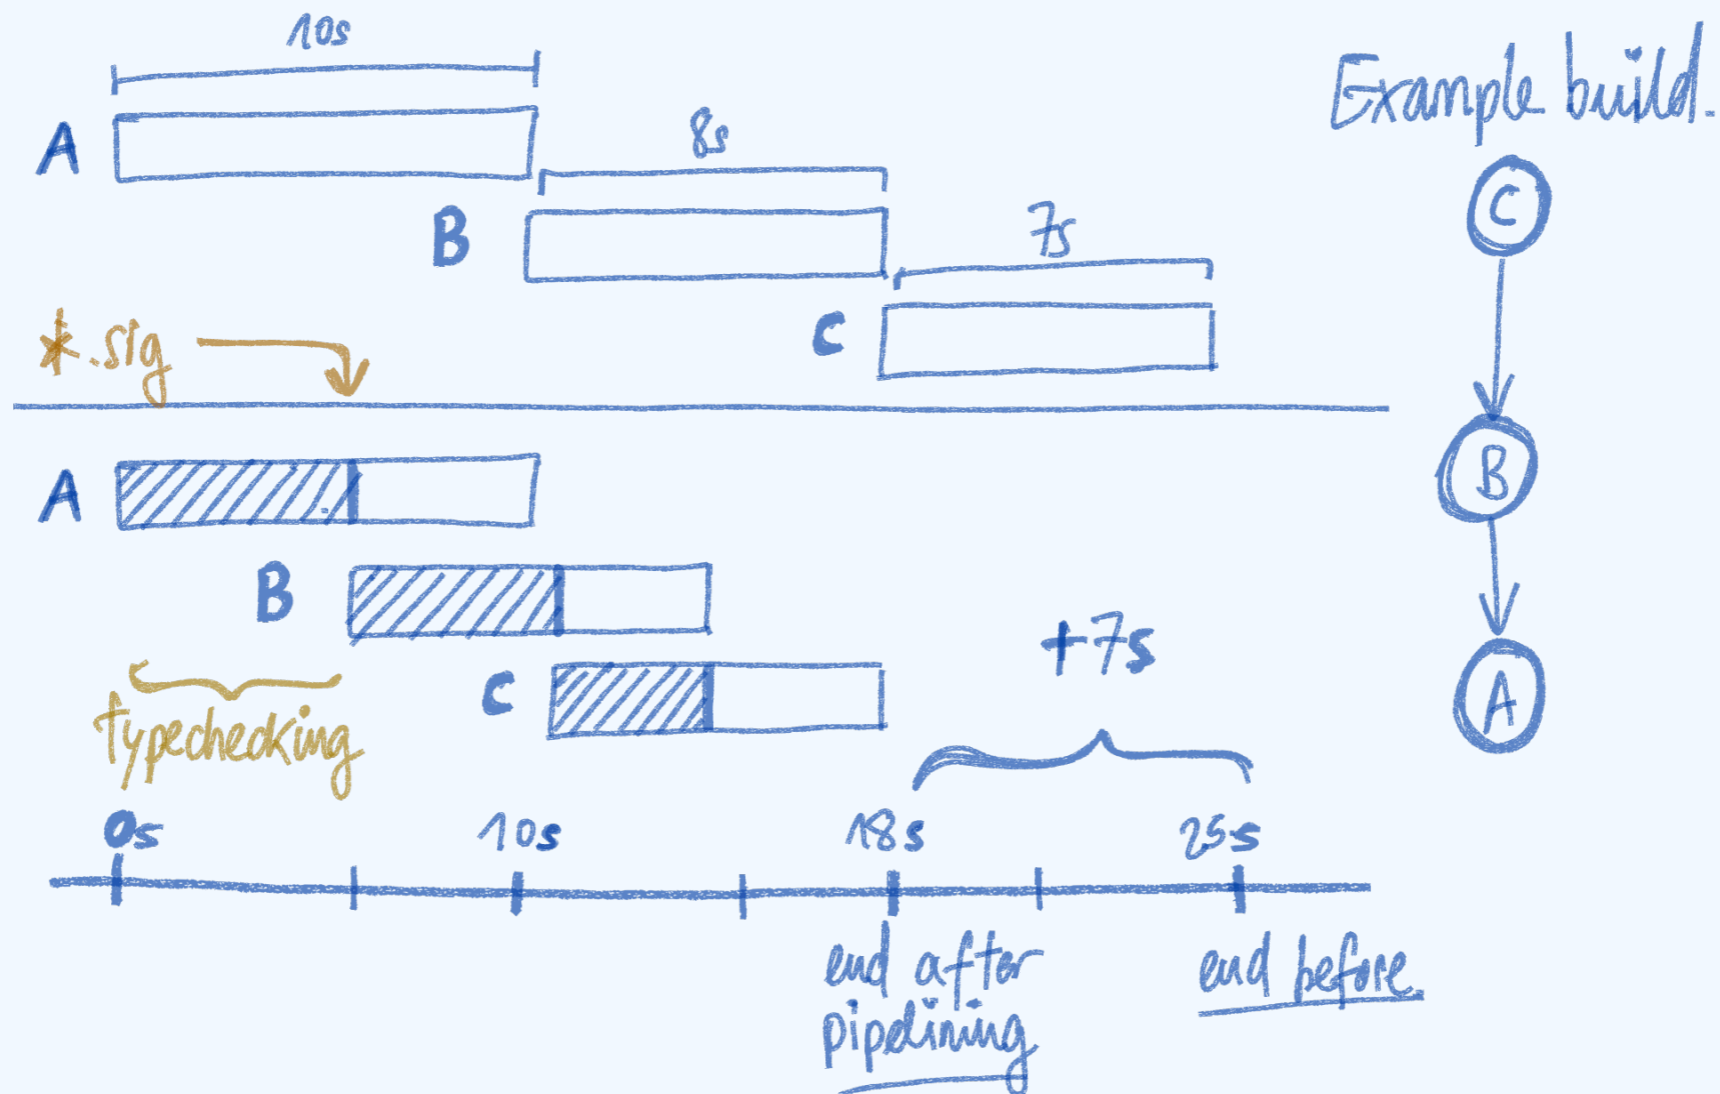


tool
maintainers

compiler
engineers

Build pipelining

A compilation technique to speed up the compilation time of a build graph by starting compilation of dependencies sooner.



Build pipelining

Upcoming blog post about it...

Example of a build-related technique that we can push directly to our users. *

Summary

Bloop is pushing the state-of-the-art of build servers to make Scala developers more productive.

Result? A tool that works & integrates with the existing ecosystem out-of-the-box.

Give it a try! scalacenter.github.io/bloop

Running, testing & Debugging
from the editor

Build pipelining
enabled by default

Future work: 1.4.0

GraalVM client binary
& shaded launcher
(no more installation)

Offloading compilation
from sbt

THE
END

Thank you!