

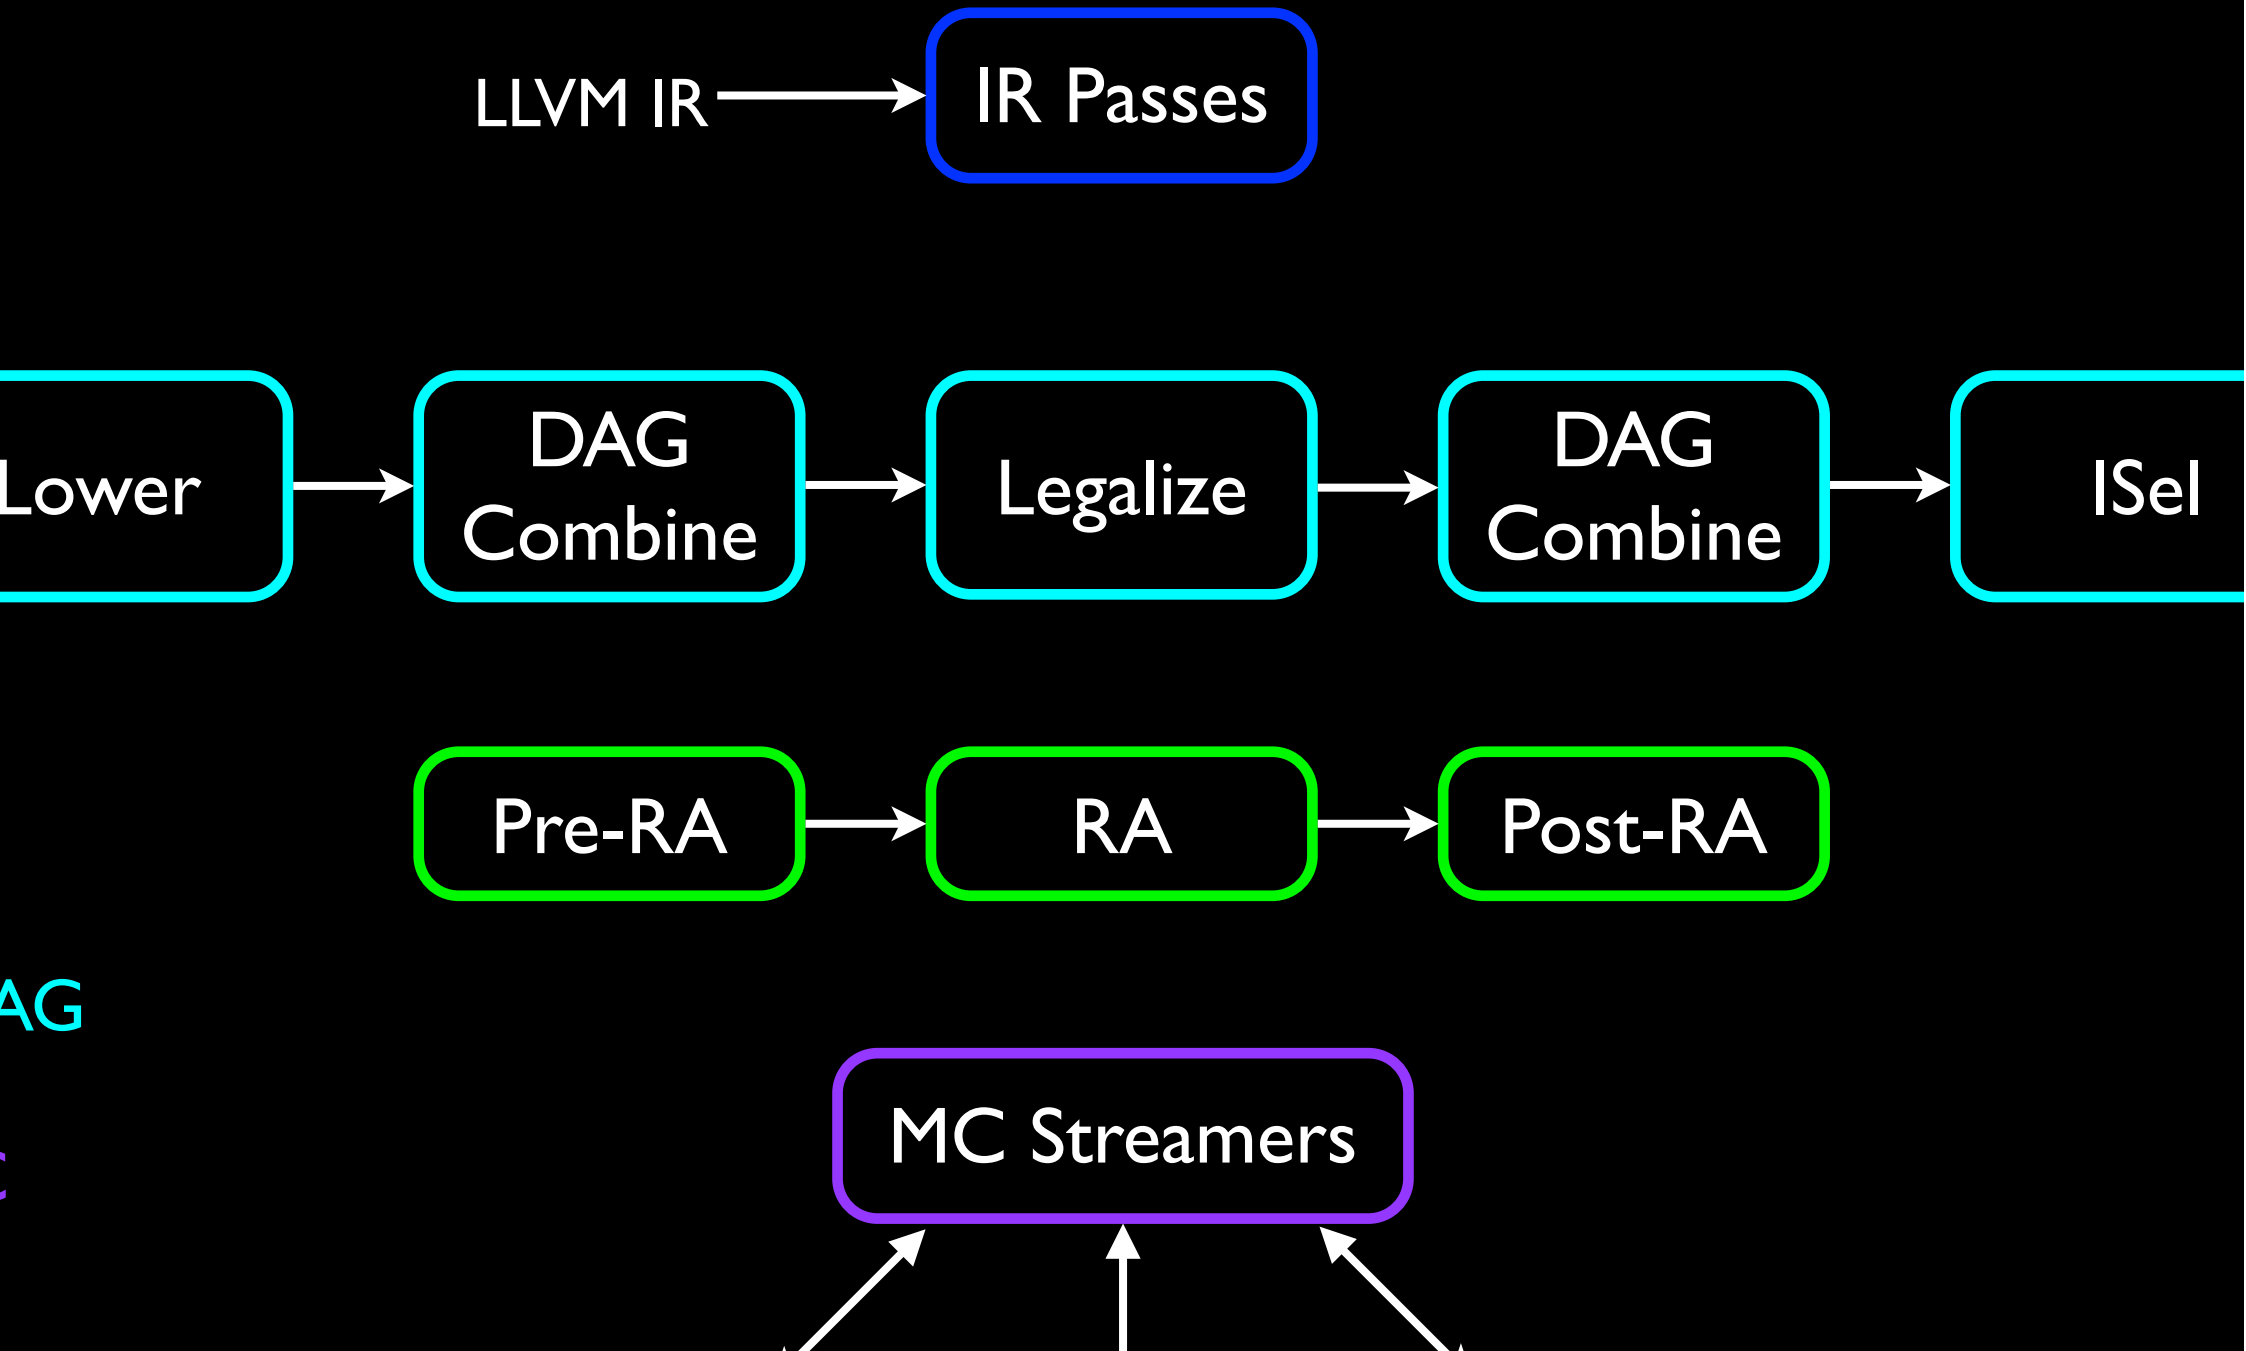
Tutorial: Building a backend in 24 hours

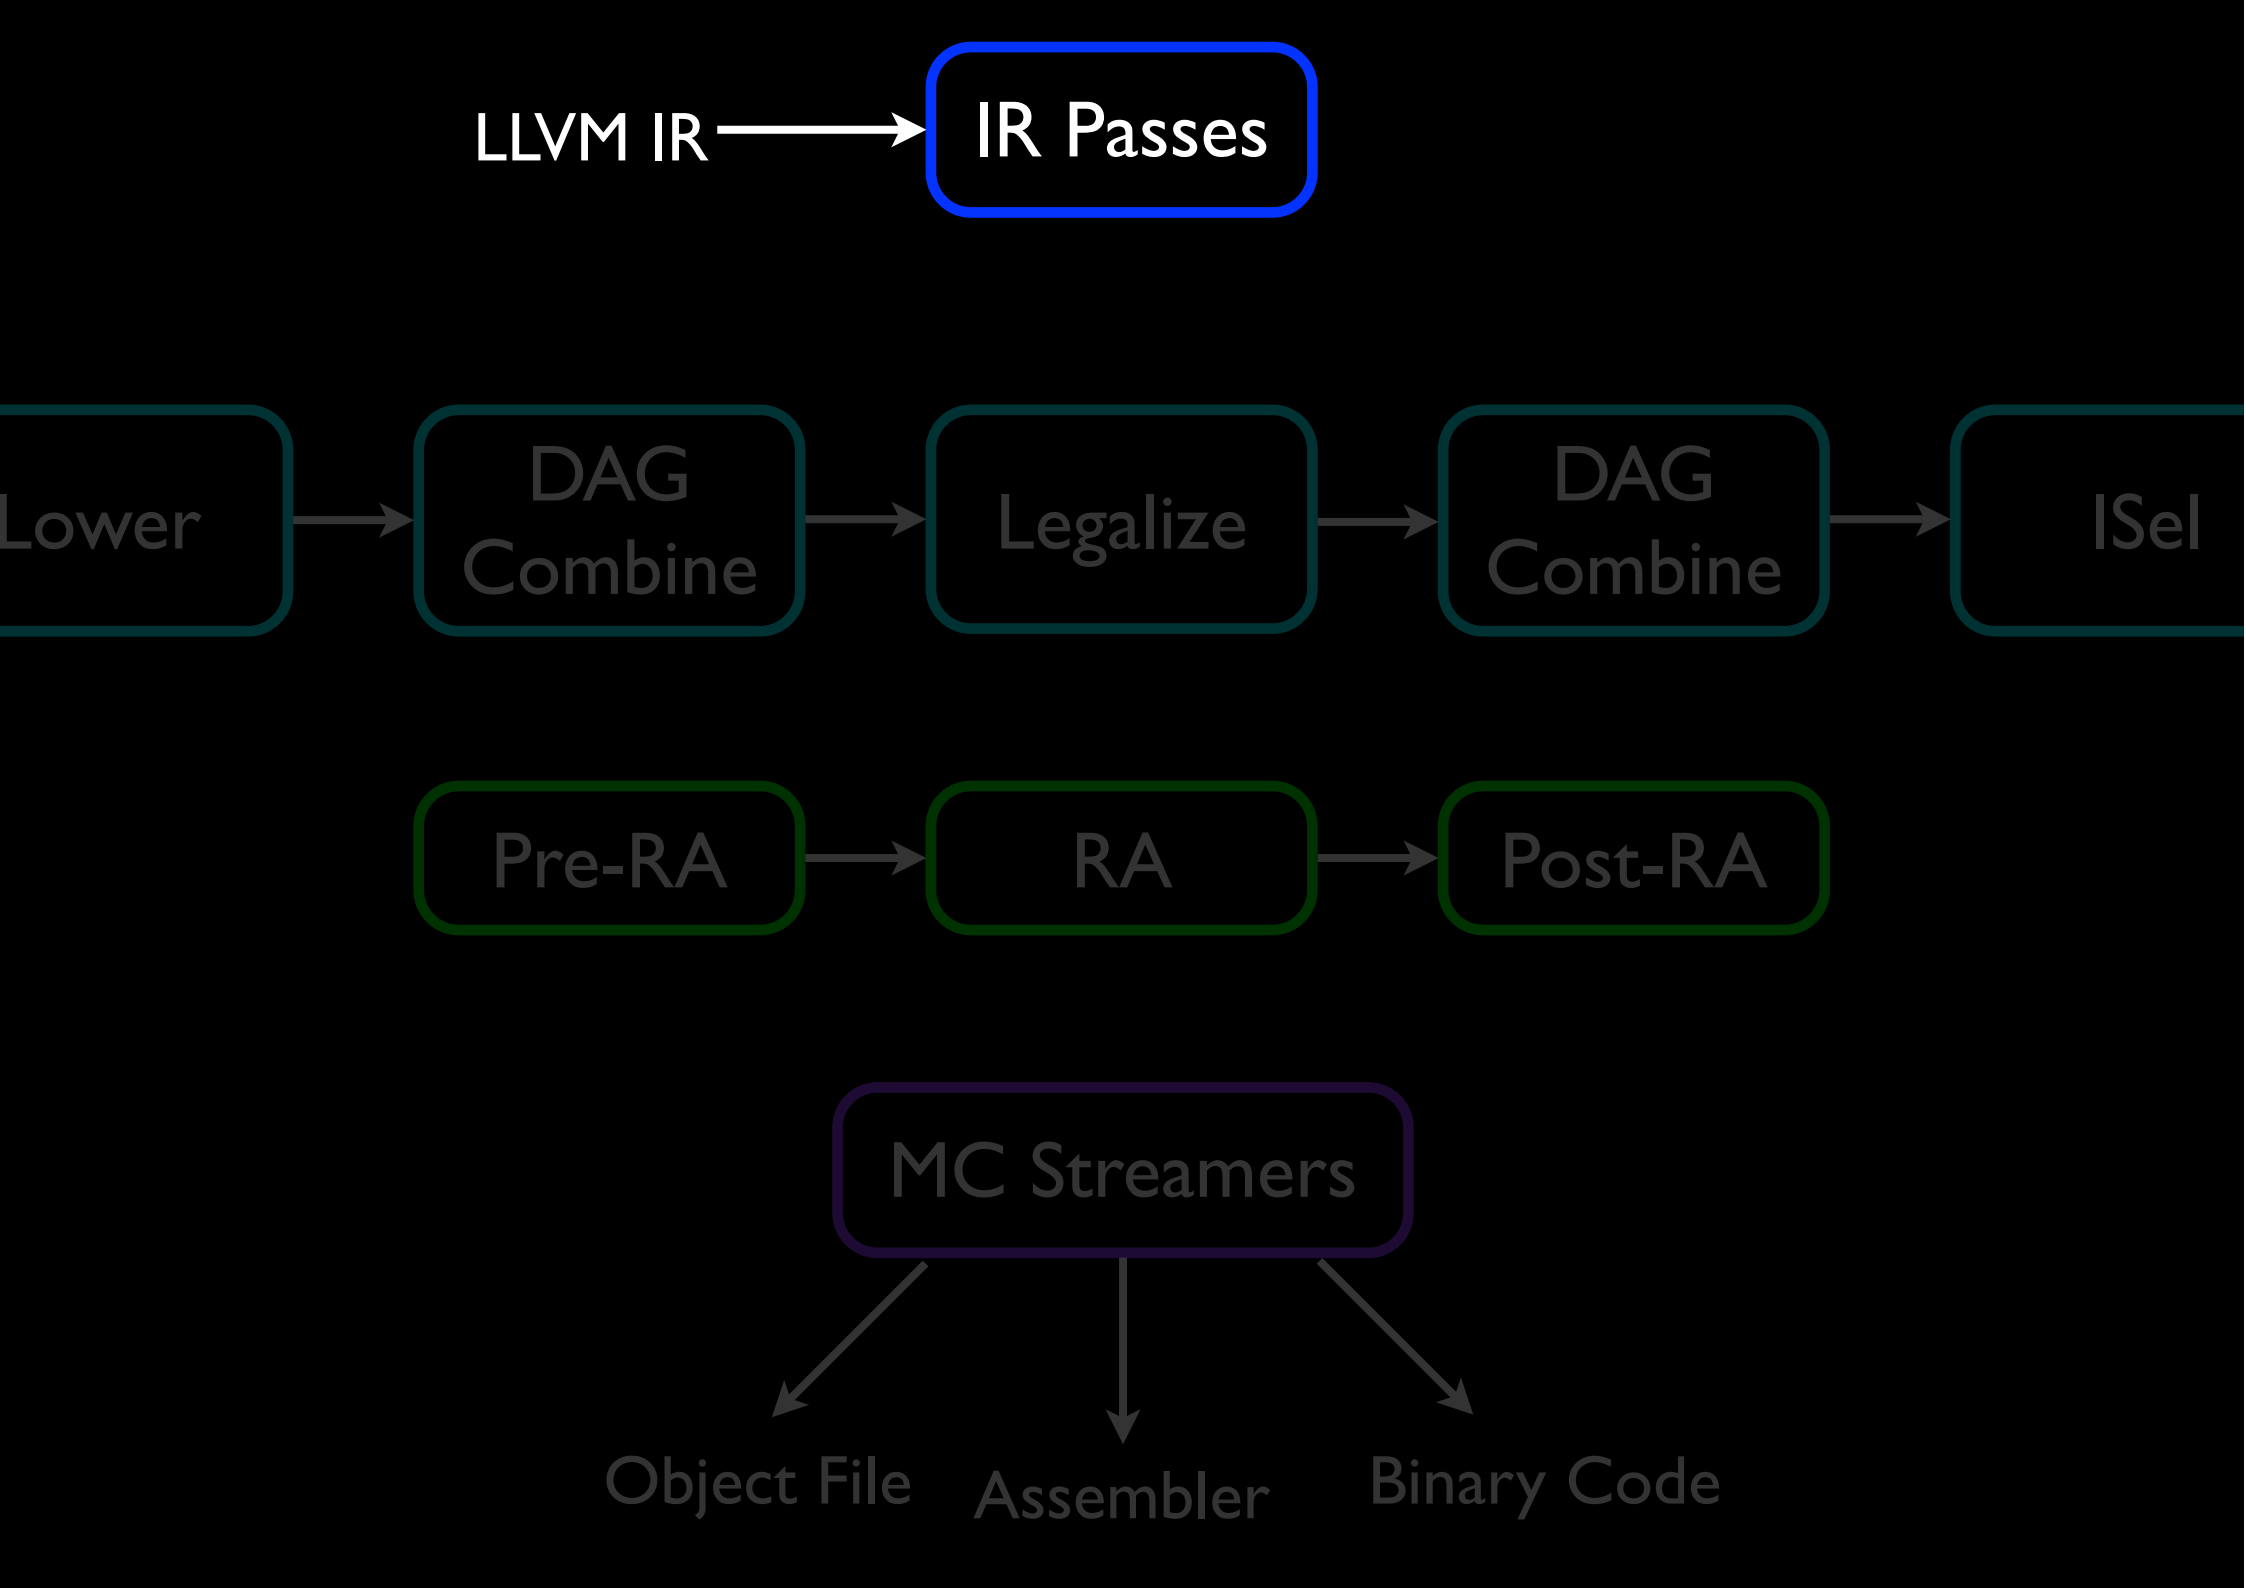
Anton Korobeynikov
anton@korobeynikov.info

Outline

1. From IR to assembler: codegen pipeline
2. MC
3. Parts of a backend
4. Example step-by-step

IR Pipeline





IR Level Passes

Why?

- Some things are easier to do at IR level
- Simplifies codegen
- Safer (pass pipeline is much more fixed)

IR Level Passes

Why?

- Some things are easier to do at IR level
- Simplifies codegen
- Safer (pass pipeline is much more fixed)

What is done?

- Late opts (LSR, elimination of dead BBs)
- IR-level lowering: GC, EH, stack protector
- Custom pre-isel passes
- CodeGenPrepare

EH Lowering

Why?

- To simplify codegen

EH Lowering

Why?

- To simplify codegen

What is done?

- Lowering of EH intrinsics to unwinding runtime constructs (e.g. `sjlj` stuff)

CodeGenPrepare

Why?

- To workaround BB-at-a-time codegen

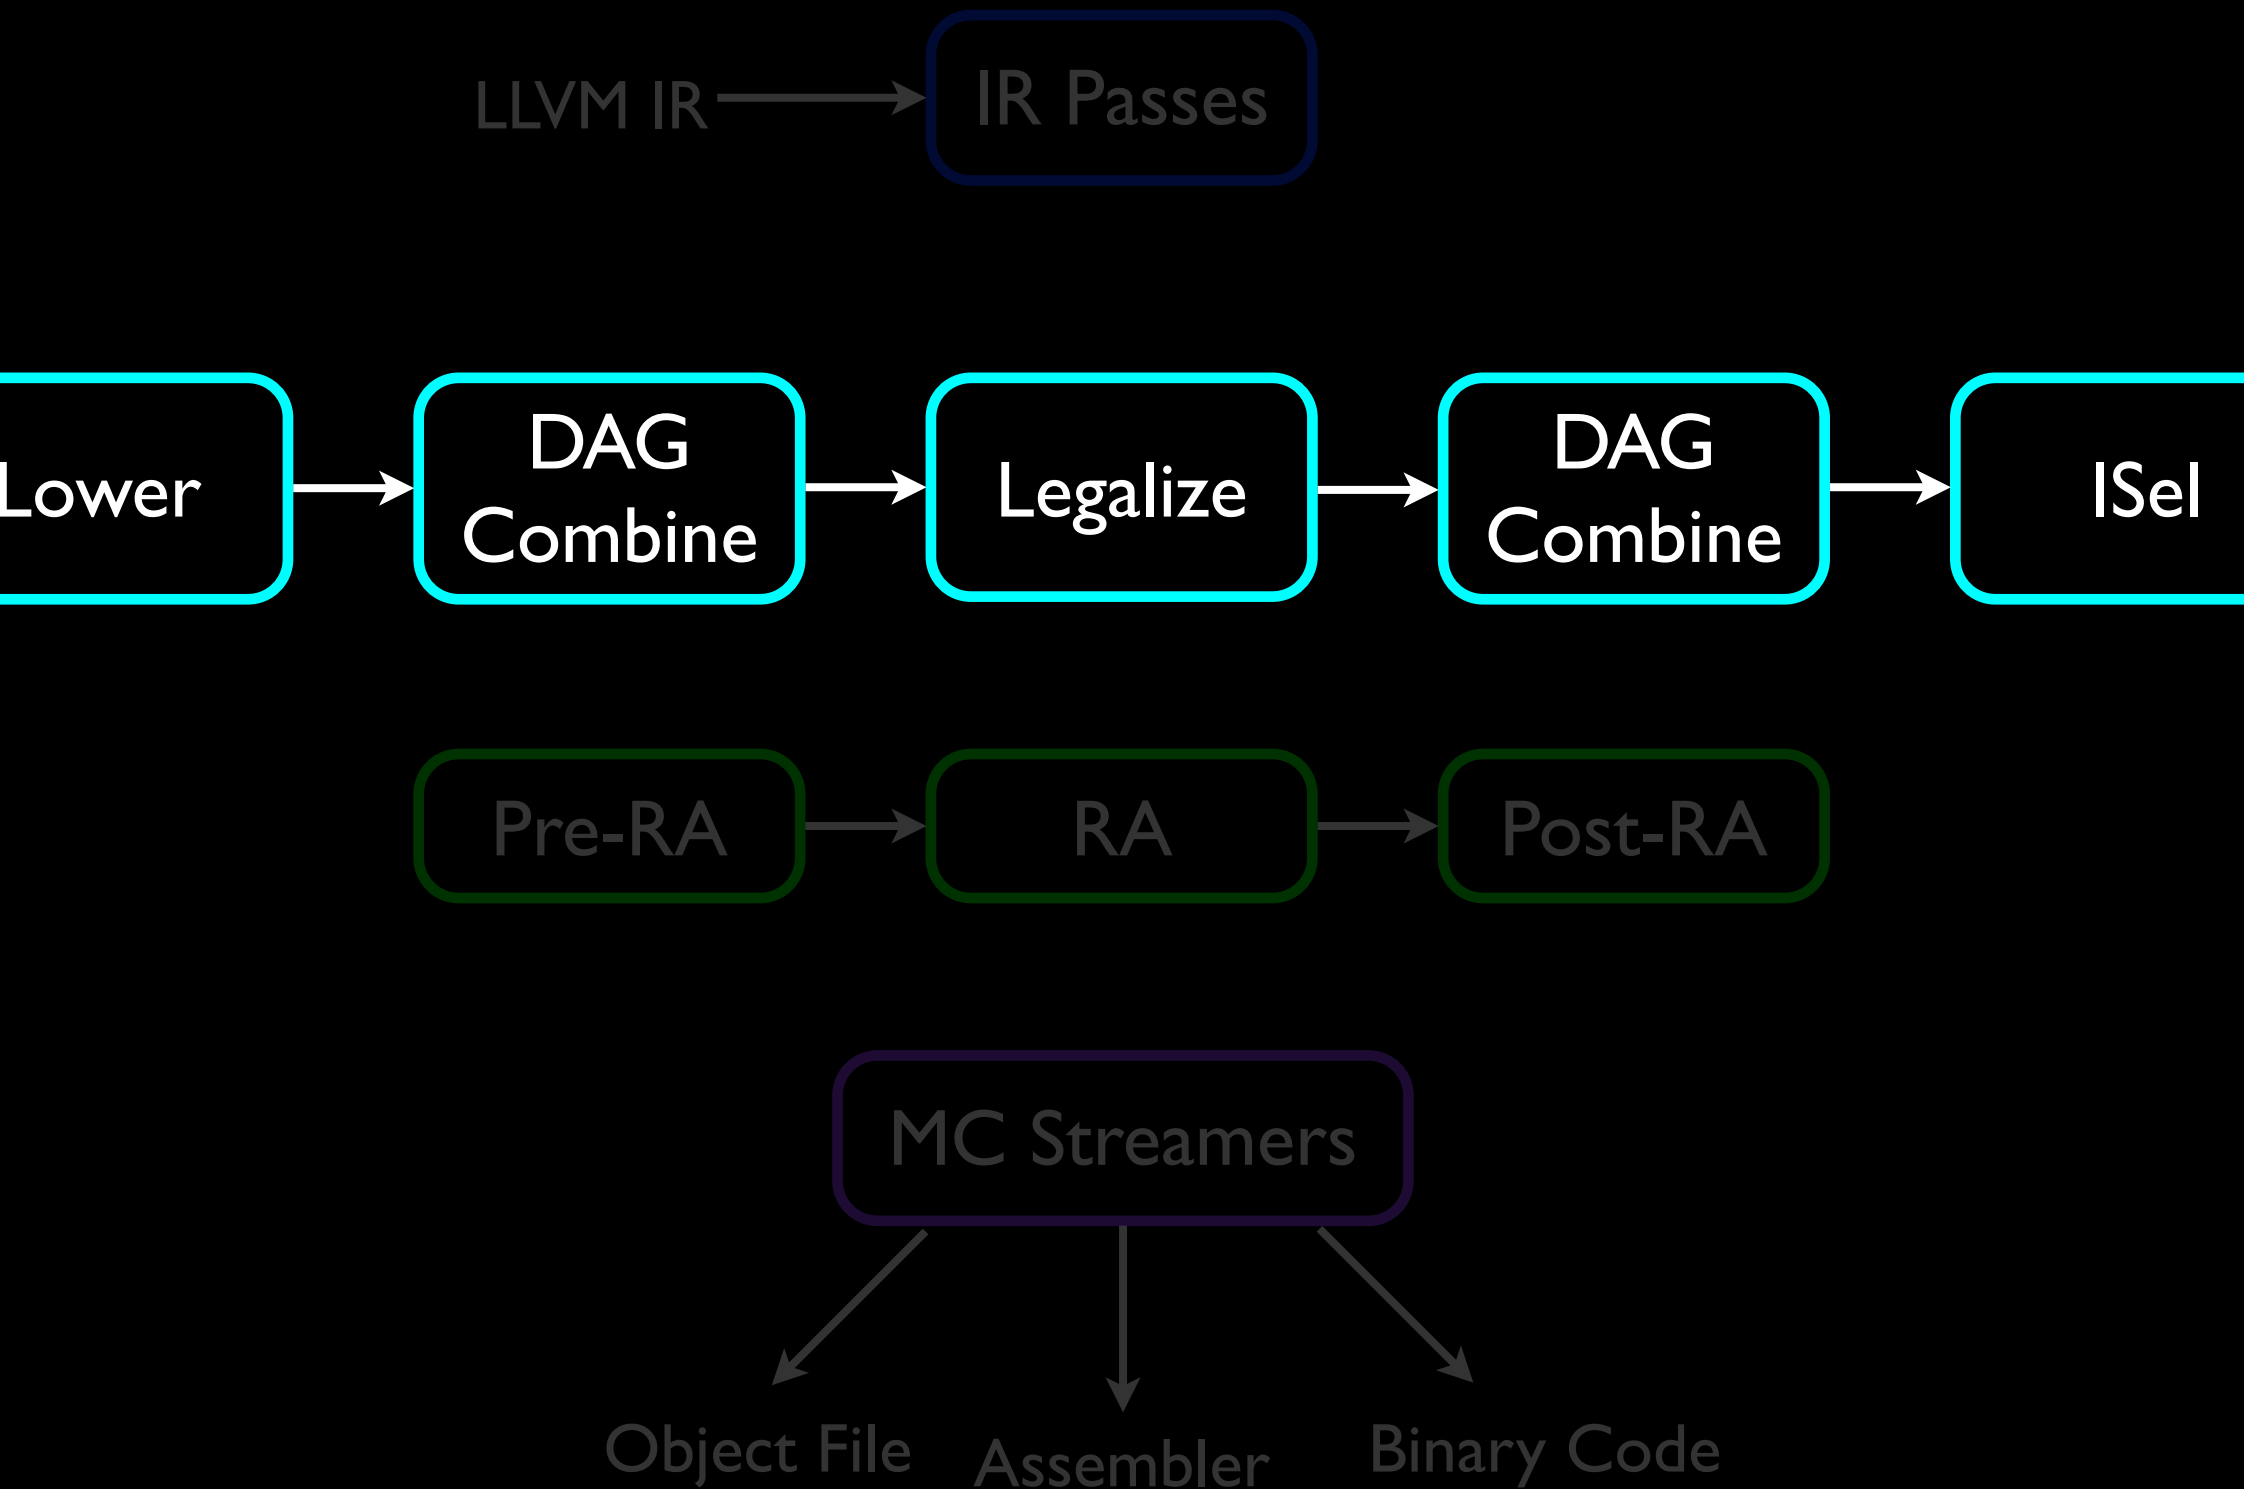
CodeGenPrepare

Why?

- To workaround BB-at-a-time codegen

What is done?

- Addressing mode-related simplifications
- Inline asm simplification (e.g. bswap patterns)
- Move debug stuff closer to defs



Selection DAG

- First strictly backend IR
- Even lower level than LLVM IR
- Use-def chains + additional stuff to keep things in order
- Built on per-BB basis

DAG-level Passes

- Lowering
- Combine
- Legalize
- Combine
- Instruction Selection

DAG Combiner

- Optimizations on DAG
- Close to target
- Runs twice - before and after legalize
- Used to cleanup / handle optimization opportunities exposed by targets

DAG Legalization

Turn non-legal operations into legal one

DAG Legalization

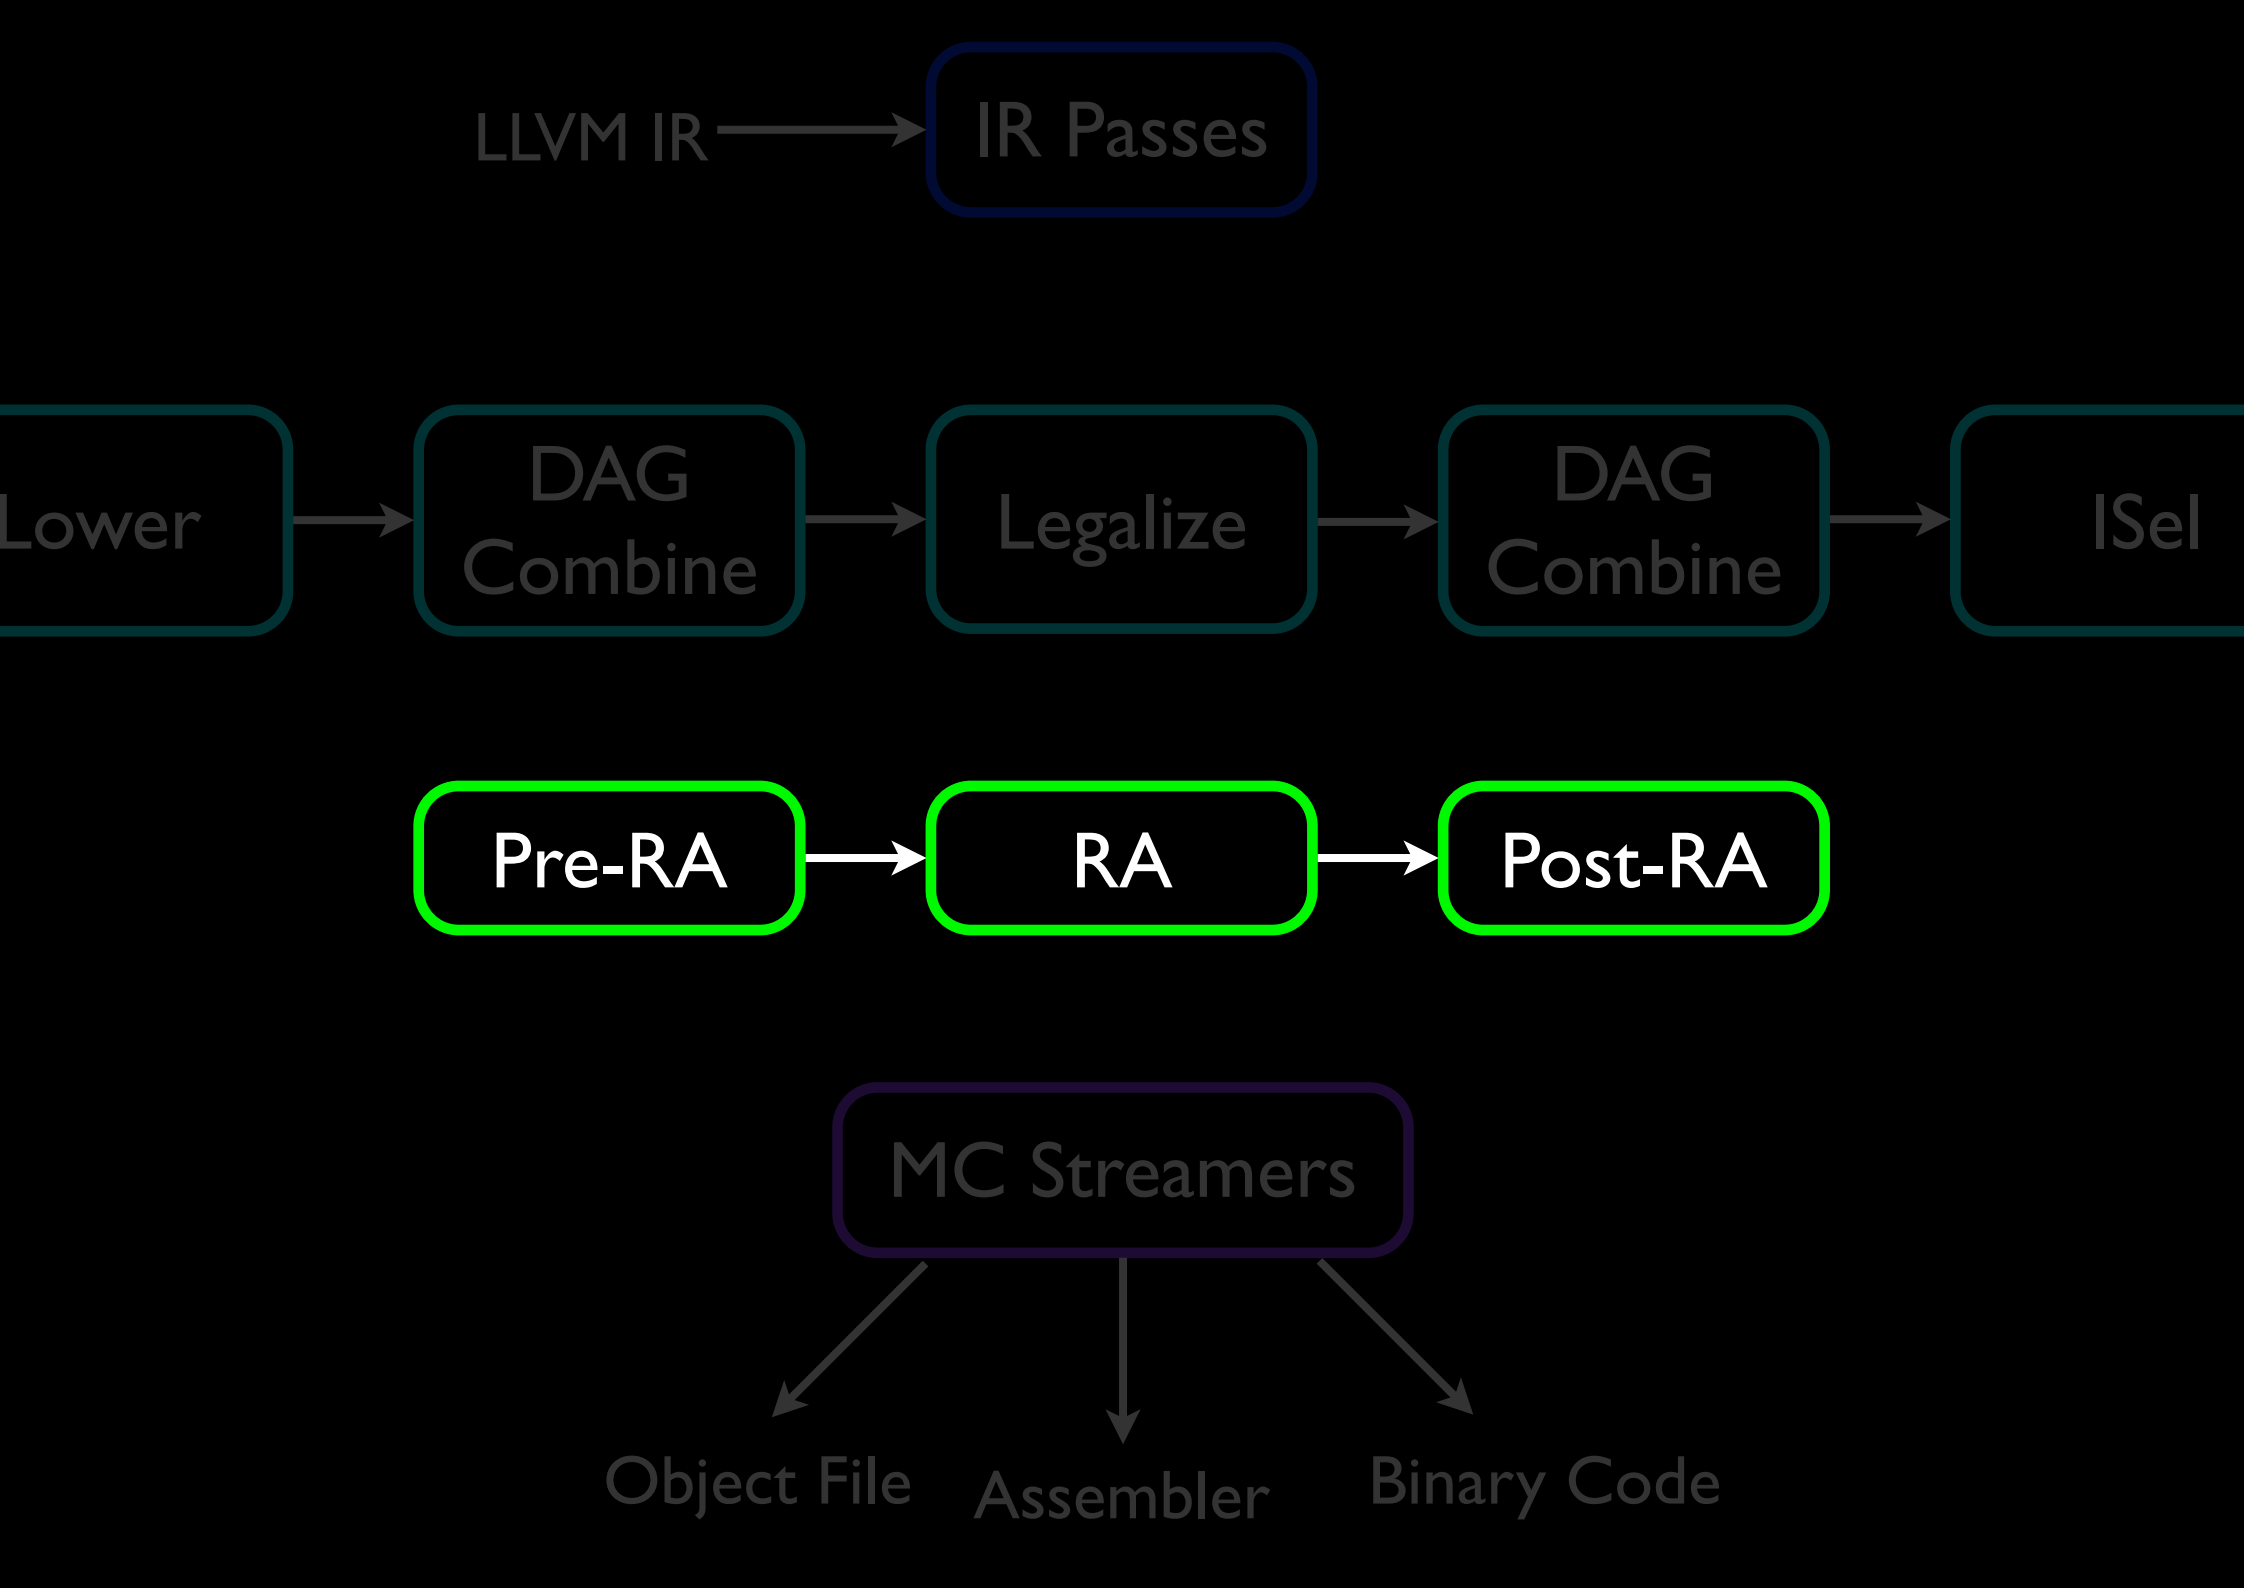
Turn non-legal operations into legal one

Examples:

- Software floating point
- Scalarization of vectors
- Widening of “funky” types (e.g. i42)

Instruction Selection

- Turns SDAGs into MIs
- Uses target-defined patterns to select instructions and operands
- Does bunch of magic and crazy pattern-matching
- Target can provide “fast but crude” isel for -O0 (fallbacks to standard one if cannot isel something)



Machine*

- Yet another set of IR: MachineInst + MachineBB + MachineFunction
- Close to target code
- Pretty explicit: set of impdef regs, basic block live in / live out regs, etc.
- Used as IR for all post-isel passes

Pre-RA Passes

- Pre-RA tail duplication
- PHI optimization
- MachineLICM, CSE, DCE
- More peephole opts

Pre-RA Passes

- Pre-RA tail duplication
- PHI optimization
- MachineLICM, CSE, DCE
- More peephole opts

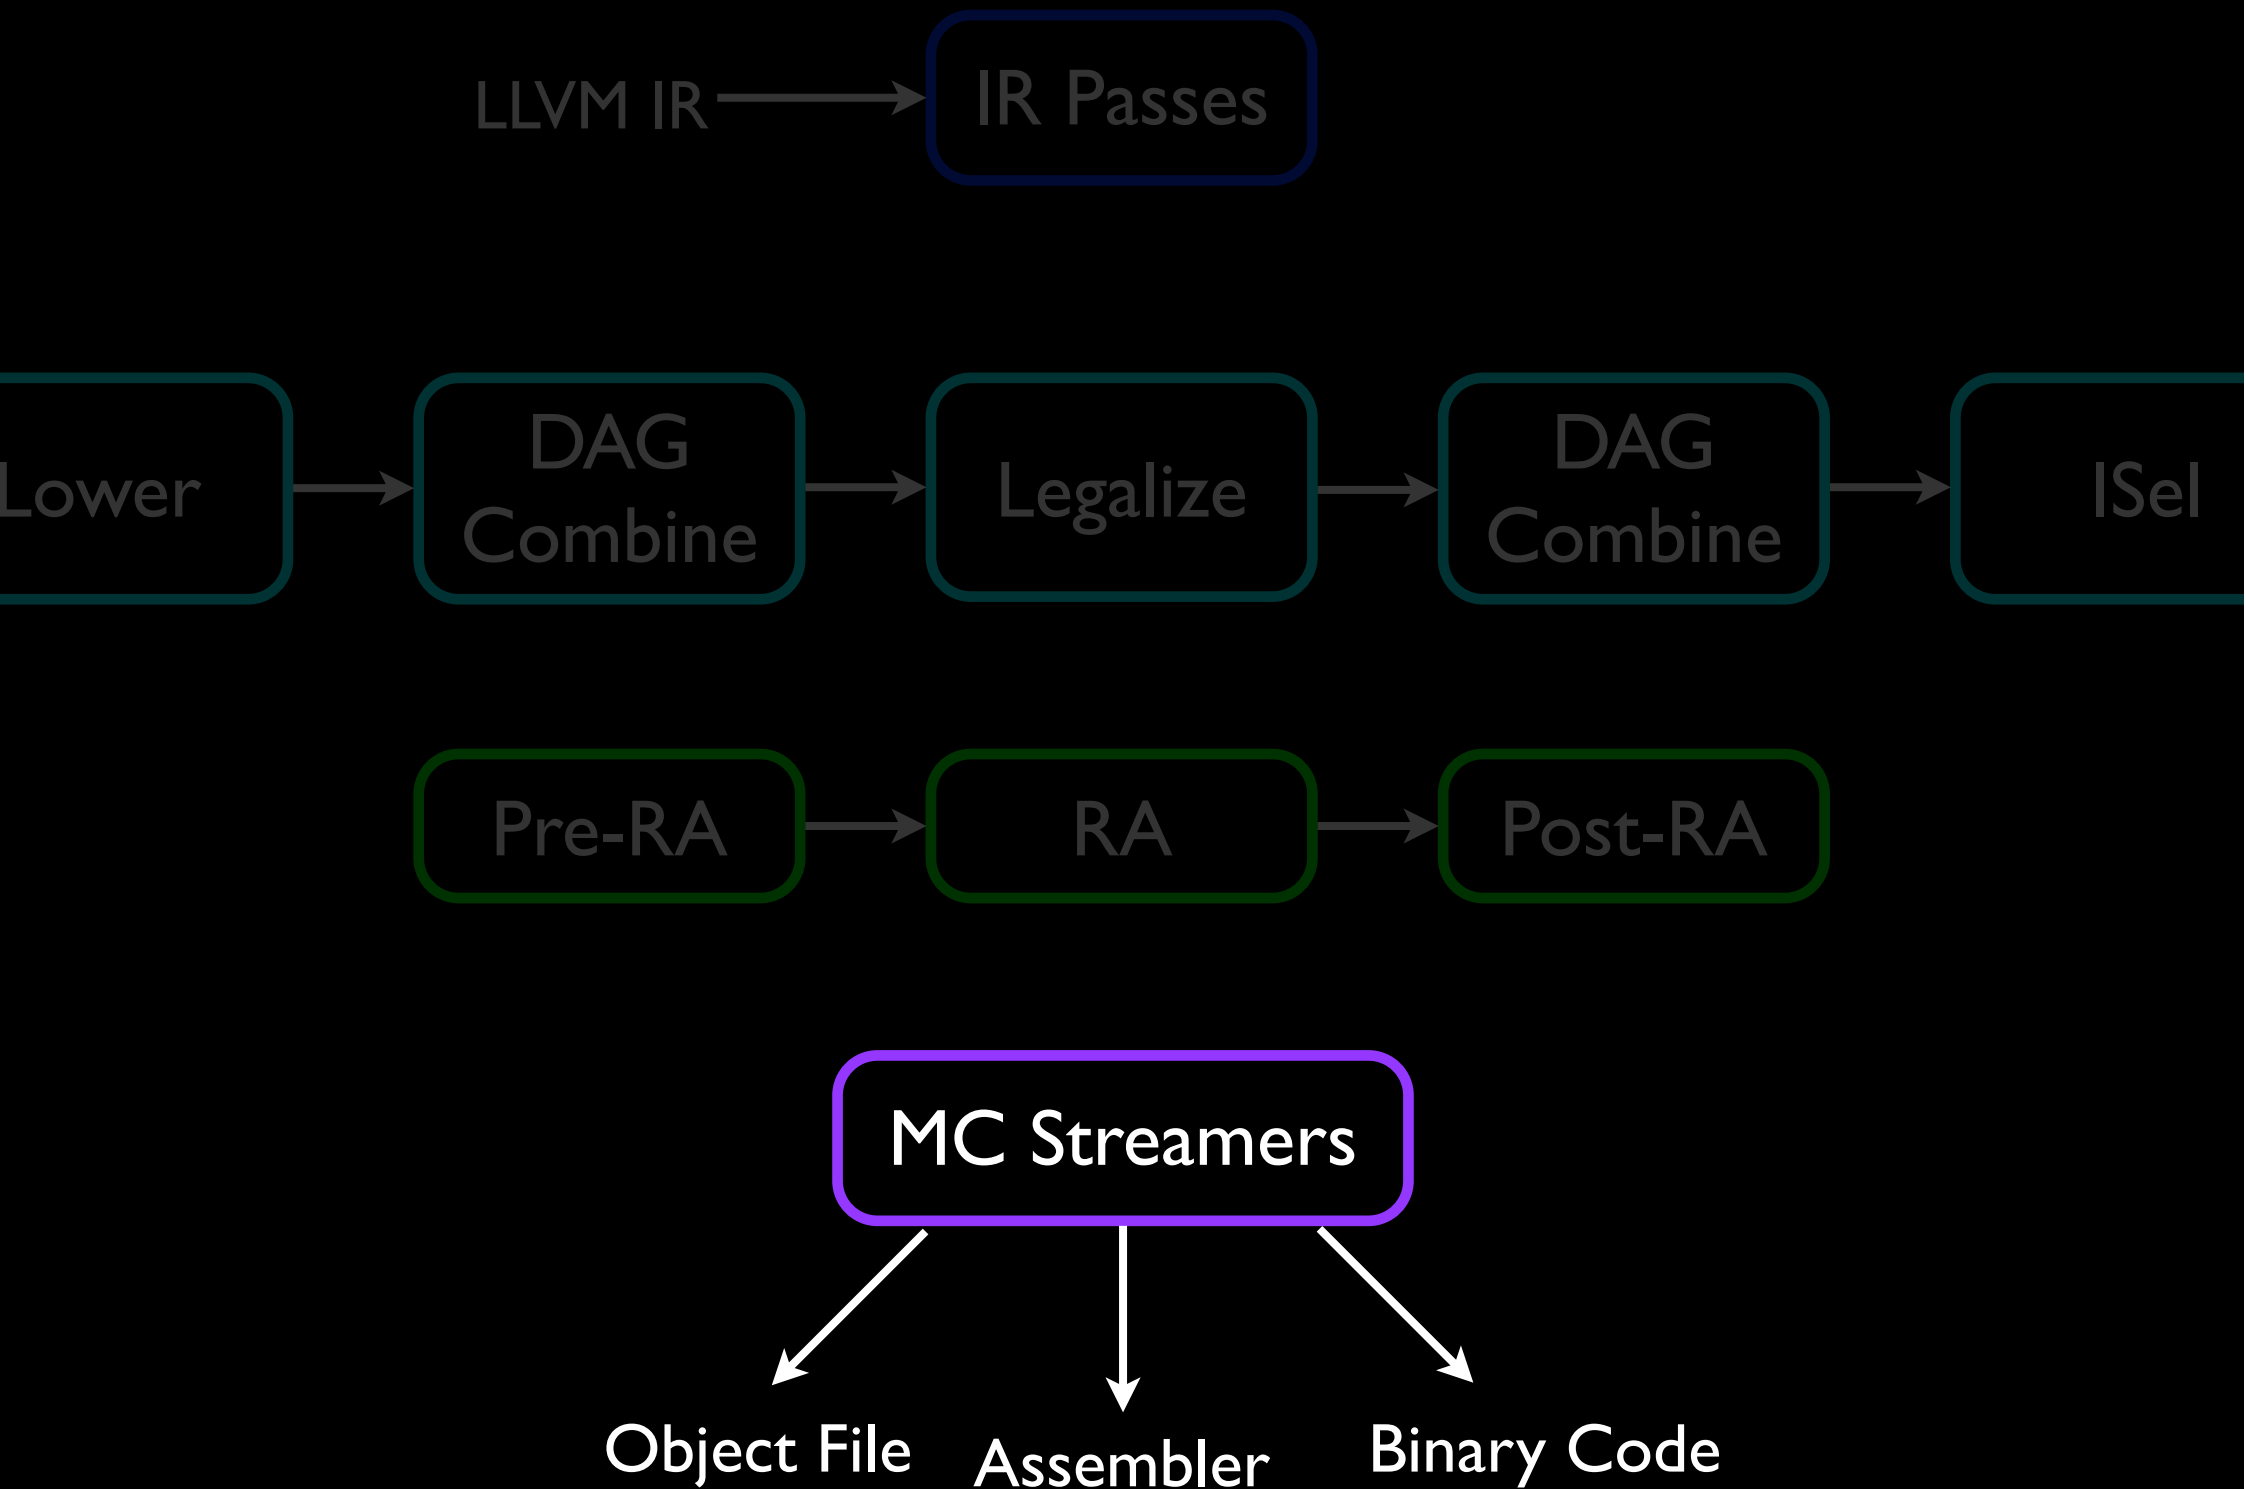
Code is still in SSA form!

Register Allocator

- Fast
- Greedy (default)
- PBQP

Post-RA Passes

1. Prologue / Epilogue Insertion & Abstract Frame Indexes Elimination
2. Branch Folding & Simplification
3. Tail duplication
4. Reg-reg copy propagation
5. Post-RA scheduler
6. BB placement to optimize hot paths



“Assembler Printing”

- Lower MI-level constructs to MCInst
- Let MCStreamer decide what to do next:
emit assembler, object file or binary code
into memory

Customization

Target can insert its own passes in specific points in the pipeline (e.g. after isel or before scheduler)

Customization

Target can insert its own passes in specific points in the pipeline (e.g. after isel or before scheduler)

Examples:

- IT block formation, load-store optimization on ARM
- Delay slot filling on MIPS or Sparc

The Backend

- Standalone library
- Mixed C++ code + TableGen
- TableGen is a special DSL used to describe register sets, calling conventions, instruction patterns, etc.
- Inheritance and overloading are used to augment necessary target bits into target-independent codegen classes

Stub Backend

How much code we need to create no-op backend?

Stub Backend

How much code we need to create no-op backend?

Some decent amount...

Stub Backend

How much code we need to create no-op backend?

Some decent amount:

- 15 classes
- around 1k LOC (both C++ and TableGen)

FooTargetMachine

- Central class in each backend
- Glues (almost) all the backend classes
- Controls the backend pipeline

FooSubtarget

- Several “subtargets” inside one target
- Usually used to model different instruction sets, platform-specific things, etc.
- Done via “subtarget features”

FooRegisterInfo

Provides various information about register sets:

1. Callee saved regs
2. Reserved (non-allocable) regs
3. Register allocation order
4. Register classes for cross-class copying & coalescing

FooRegisterInfo

Provides various information about register sets:

1. Callee saved regs
2. Reserved (non-allocable) regs
3. Register allocation order
4. Register classes for cross-class copying & coalescing

FooRegisterInfo.td

TableGen description of:

1. Registers,
2. Sub-registers (and aliasing sets for regs)
3. Register classes

FoolSelLowering

- Central class for target-aware lowering
- Turns target-neutral SelectionDAG in target-aware (suitable for instruction selection)
- Something can be lowered (albeit not efficiently) in generic way
- Some cases (e.g. argument lowering) always require custom lowering

FooCallingConv.td

Describes the calling convention:

1. What & where & in which order should be passed
2. Not self-containing: used to simplify custom lowering routines
3. Autogenerate set of callee-save registers

FoolSelDAGToDAG

- Does most of instruction selection
- Most of C++ code is autogenerated from instruction patterns
- Custom instruction selection code:
 - Complex addressing modes
 - Instructions which require additional care

FoolInstrInfo

Hooks used by codegen to:

1. Emit reg-reg copies
2. Save / restore values on stack
3. Branch-related operations
4. Determine instruction sizes

FoolnstrInfo.td

Defines the instruction patterns:

- DAG: level of input & output operands
- MI: Instruction Encoding
- ASM: Assembler printing strings

FoolInstrInfo.td

Defines the instruction patterns:

- DAG: level of input & output operands
- MI: Instruction Encoding
- ASM: Assembler printing strings

TableGen magic can autogenerate many things