

Generate Compilers from Hardware Models!

Gus Smith, PhD Candidate, University of Washington
PLARCH 2023




Hi everybody. My name is Gus Smith, and I'm a PhD candidate at the University of Washington's PLSE lab. Today I'm going to be talking about why we should generate compilers from hardware models.



Ben Kushigian



Vishal Canumalla



Andrew Cheung



René Just



Zachary Tatlock

First, let me shout out my coauthors in PLSE: Ben, Vishal, Andrew, Rene, and Zach, some of whom are in the audience today!

The Hardware Lottery

Sara Hooker

Google Research, Brain Team

shooker@google.com

2020

In 2020, Sara Hooker from Google Research released a paper titled the Hardware Lottery. The thesis of this paper is that

[build] ...

She argues that, for example, the current track of matrix-multiplication-based advances in machine learning are undeniably linked to the abundance of hardware for matrix multiplication, and that other research directions in machine learning are subsequently less likely to be successful.

The hardware lottery is a direct challenge to us in this room. By my reading, the takeaway for our community is that

[build] new platforms (that is, new hardware and their associated compiler stack) should be easier to build, so that the best ideas win — not just the ideas with hardware on their side!

In this talk, I'm going to focus on what I know:

[build] compilers.

So, let's ask the question,

[build] why *are* compilers so difficult

The Hardware Lottery

Sara Hooker

Google Research, Brain Team

shooker@google.com

2020

Hardware and compilers have a disproportionate role in deciding which research ideas succeed or fail.

The Hardware Lottery

Sara Hooker

Google Research, Brain Team

shooker@google.com

2020

Hardware and compilers have a disproportionate role in deciding which research ideas succeed or fail.

Takeaway for our community: new platforms (hardware + compiler stack) should be easier to build, so that the best ideas win!

The Hardware Lottery

Sara Hooker

Google Research, Brain Team

shooker@google.com

2020

Hardware and compilers have a disproportionate role in deciding which research ideas succeed or fail.

Takeaway for our community: new platforms (hardware + compiler stack) should be easier to build, so that the best ideas win!

The Hardware Lottery

Sara Hooker

Google Research, Brain Team

shooker@google.com

2020

Hardware and compilers have a disproportionate role in deciding which research ideas succeed or fail.

Takeaway for our community: new platforms (hardware + compiler stack) should be easier to build, so that the best ideas win!

Why are compilers hard to build?

Why are compilers hard to build?

Imagine we have

[build] a hardware engineer who

[build] builds a new hardware design. If she wants to run

[build] programs on her hardware, she'll need a

[build] compiler that

[build] compiles programs to run on her design. This requires

[build] a compiler engineer, who, with much time and effort

[build] builds a compiler, whose design is informed not only by

[build] the hardware design itself, but also by

[build] communications with the hardware designer,

[build] any documentation that exists for the design, and finally, by

[build] the compiler engineer's own internal model of how the hardware works. If that's not confusing enough, if the team wants to ensure the compiler is correct, they'll hire

[build] a verification engineer to

[build] build a

[build] formal verification model, which

[build] models the hardware and

[build] verifies that the compiler is correct. Similarly to the compiler, the design of the verification model is informed by

[build] communications with the hardware designer,


[build] the hardware design itself,

[build] the documentation, and
[build] the verification engineer's own ideas about how the hardware works.
[build] If this is looking a little complicated, well, I agree!
But the process of building a compiler is more than just confusing; it also
[build] requires significant developer effort,

Why are compilers hard to build?



Why are compilers hard to build?



Why are compilers hard to build?



builds →

Hardware
Design

Why are compilers hard to build?



Programs


Compiler



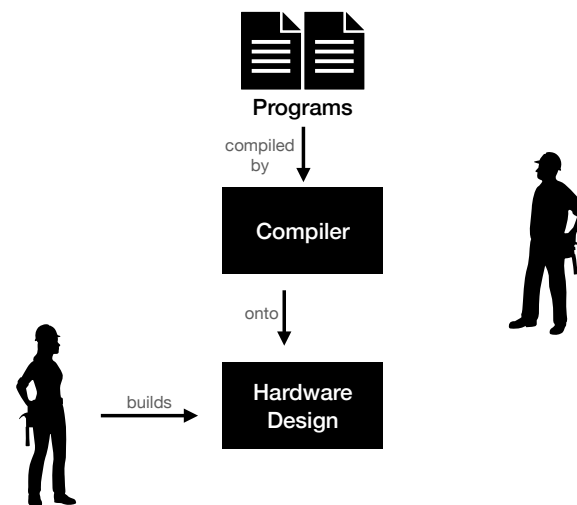
builds →

Hardware
Design

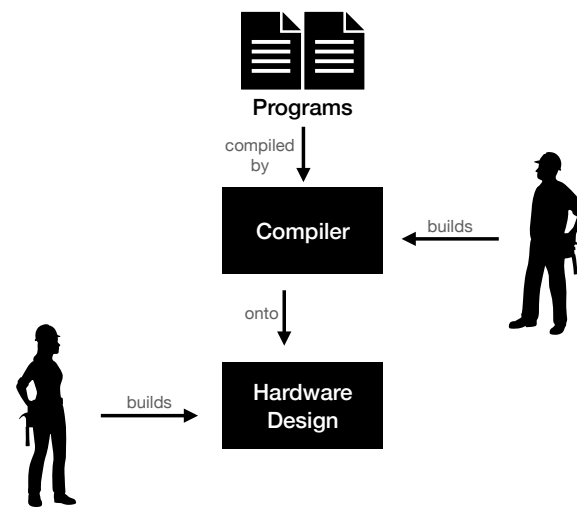
Why are compilers hard to build?



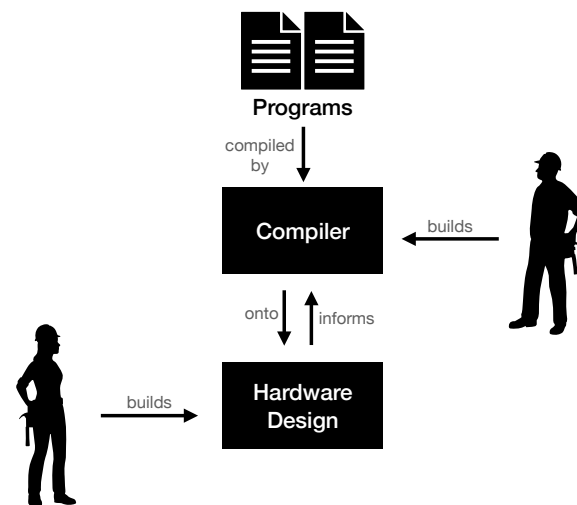
Why are compilers hard to build?



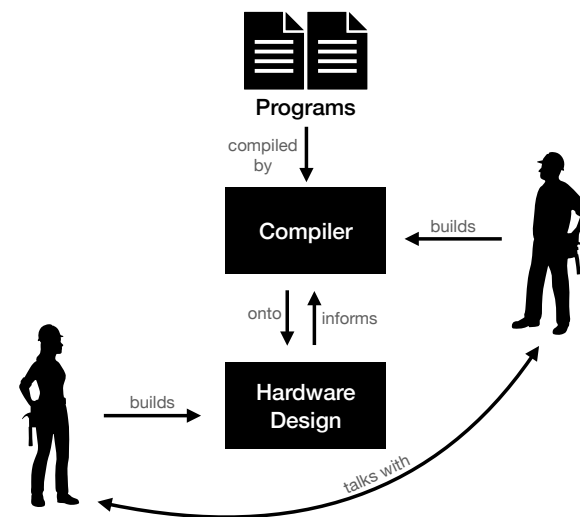
Why are compilers hard to build?



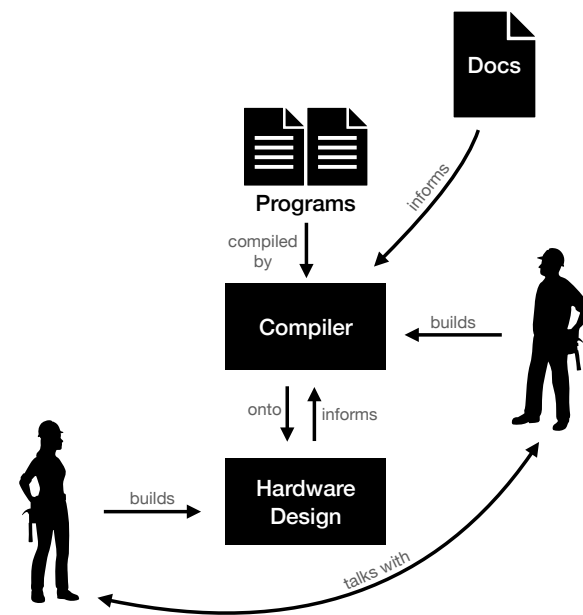
Why are compilers hard to build?



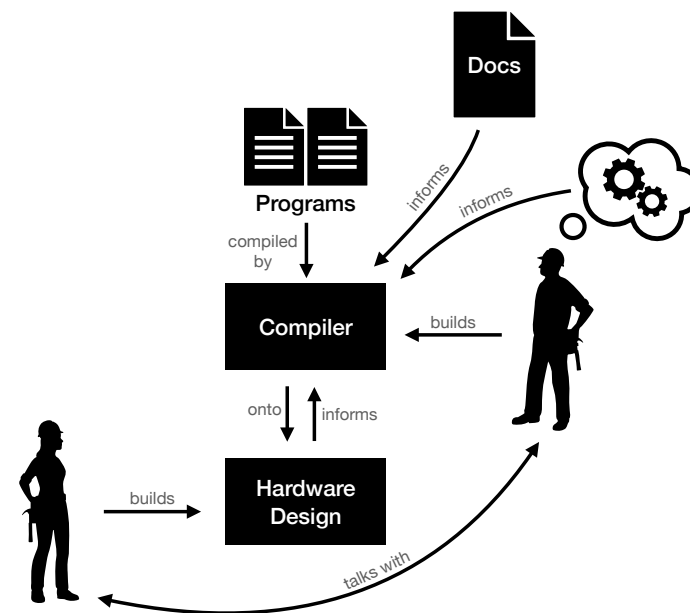
Why are compilers hard to build?



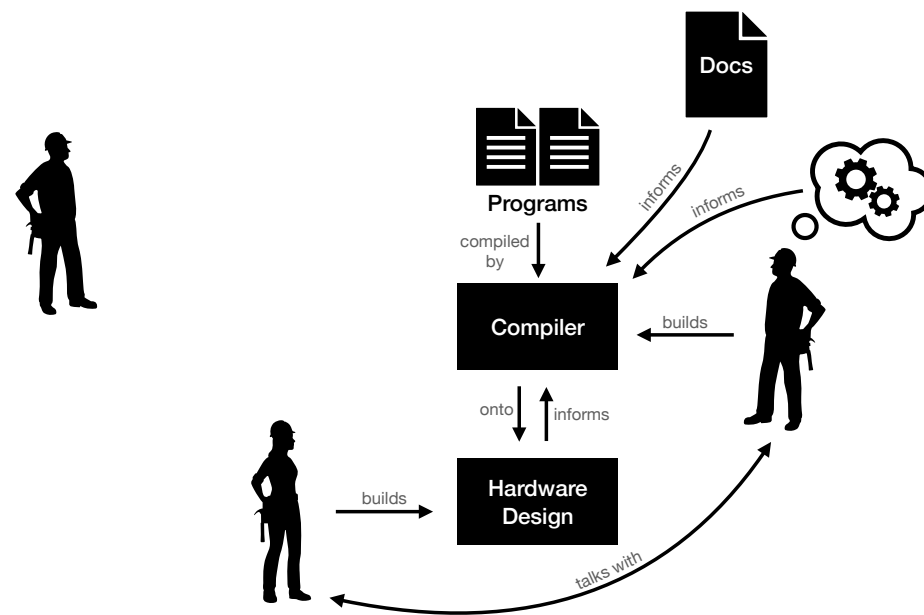
Why are compilers hard to build?



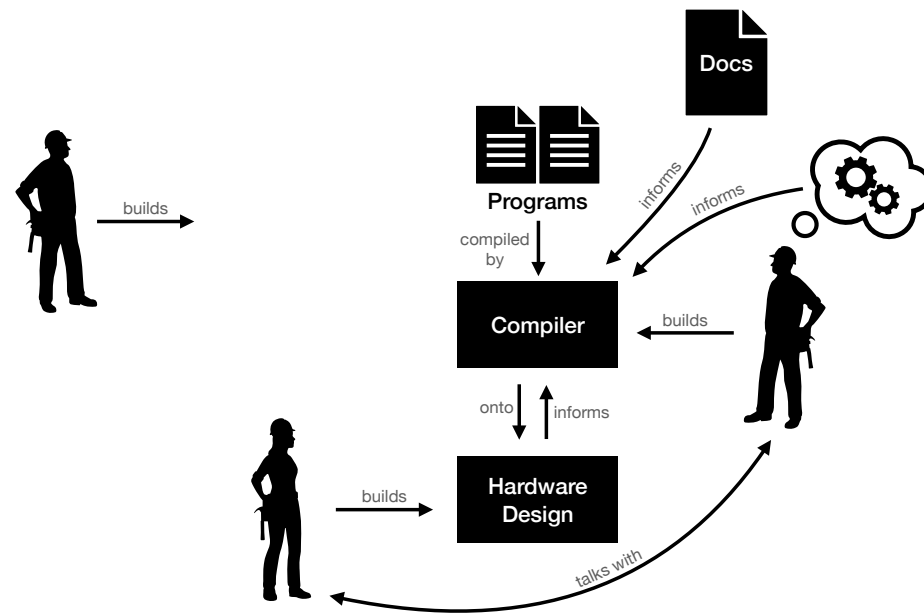
Why are compilers hard to build?



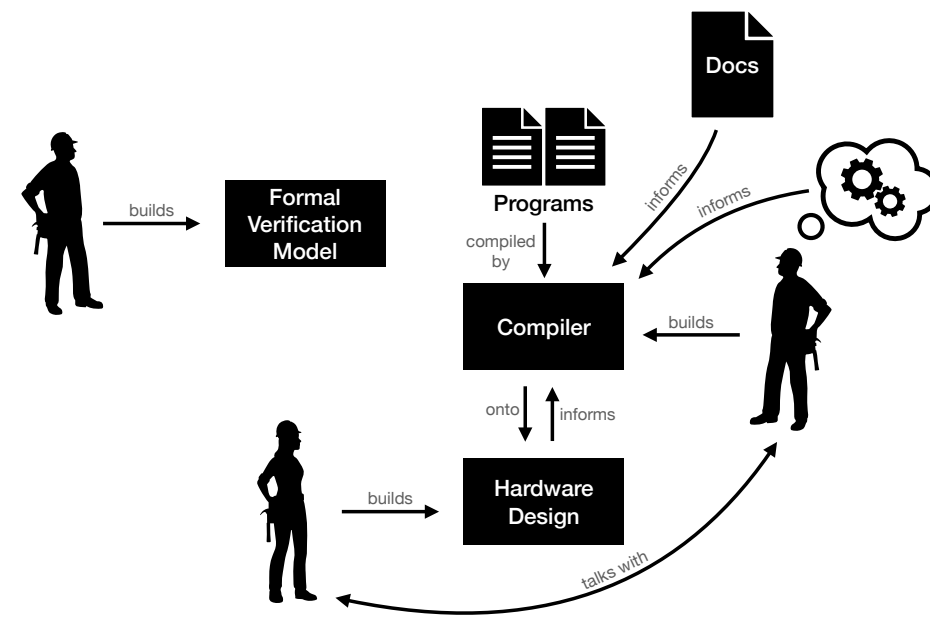
Why are compilers hard to build?



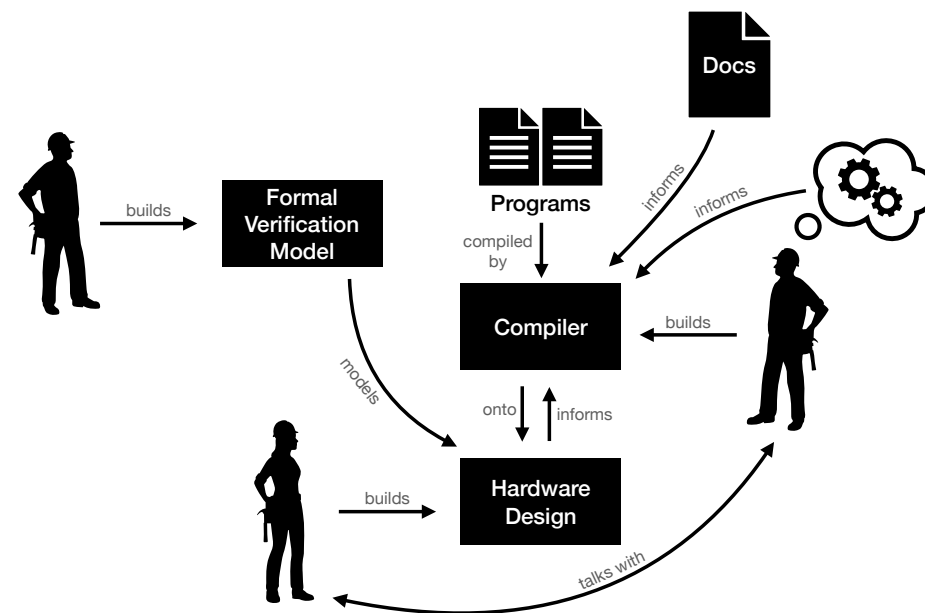
Why are compilers hard to build?



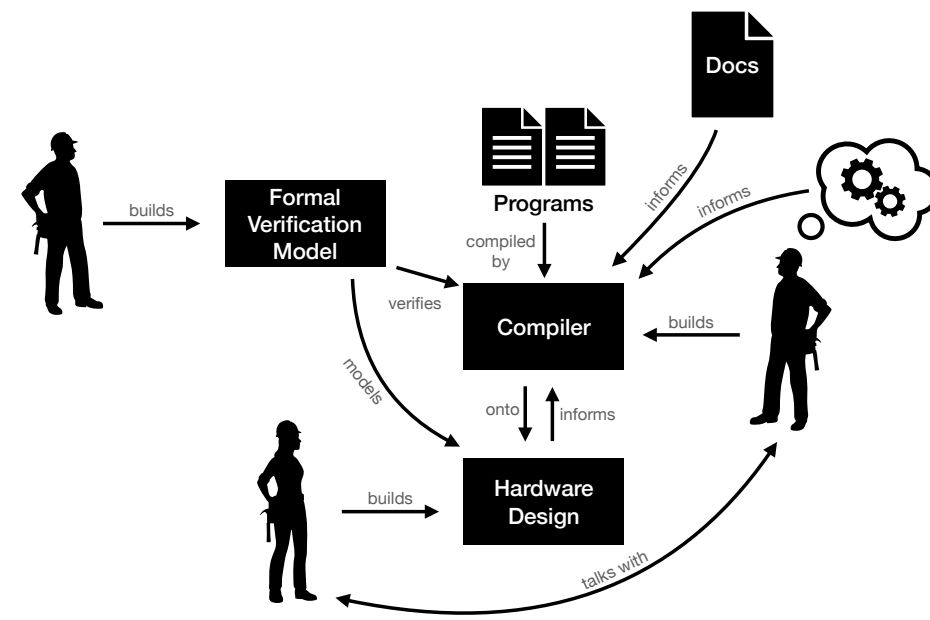
Why are compilers hard to build?



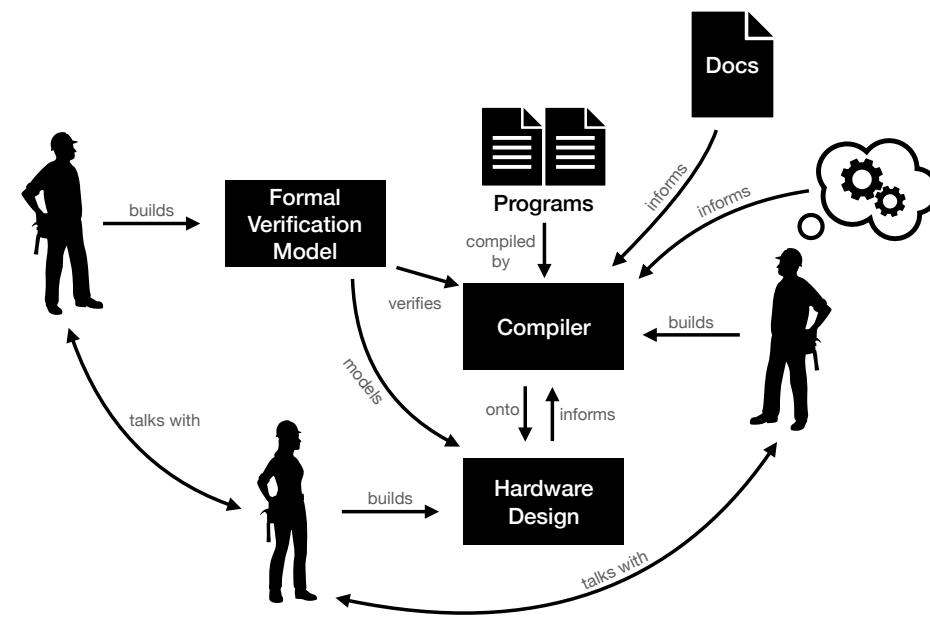
Why are compilers hard to build?



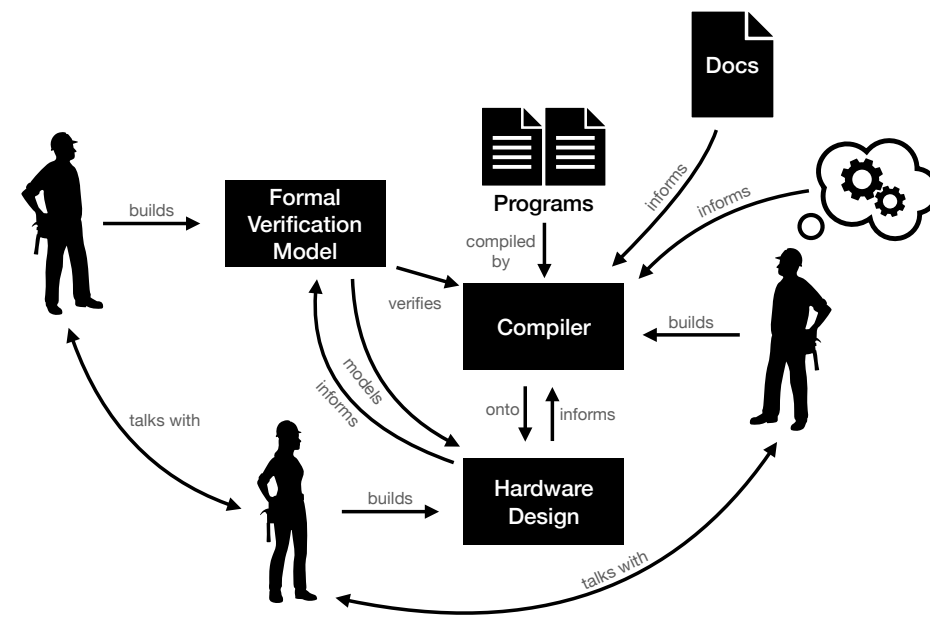
Why are compilers hard to build?



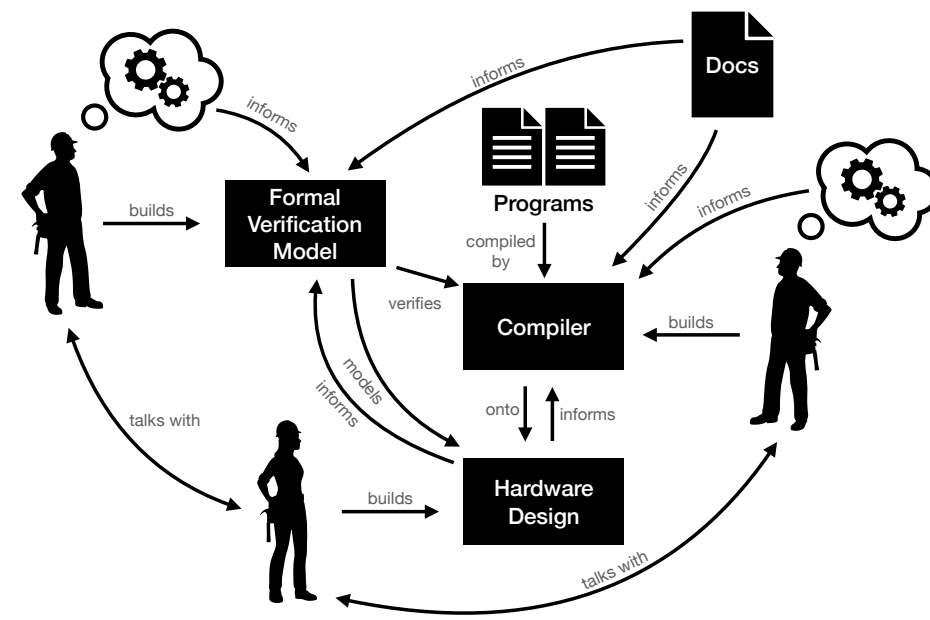
Why are compilers hard to build?

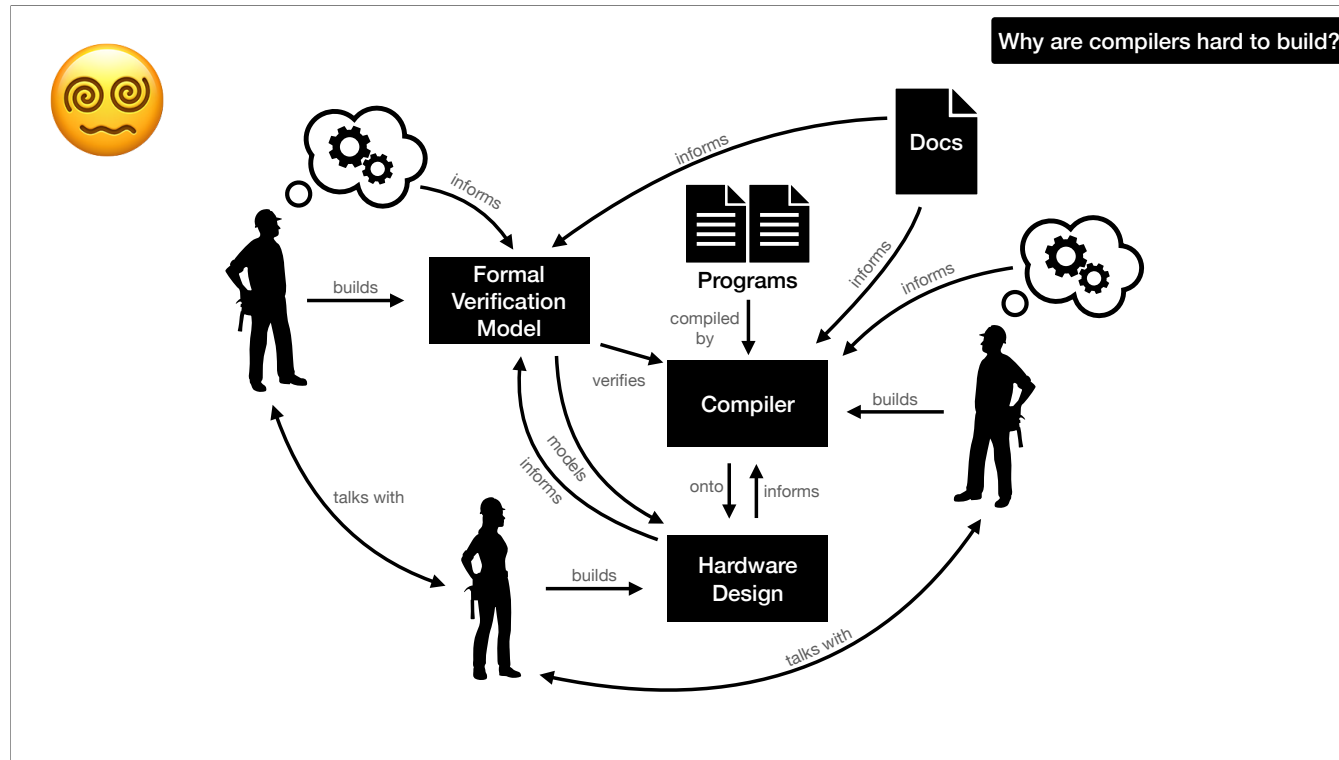


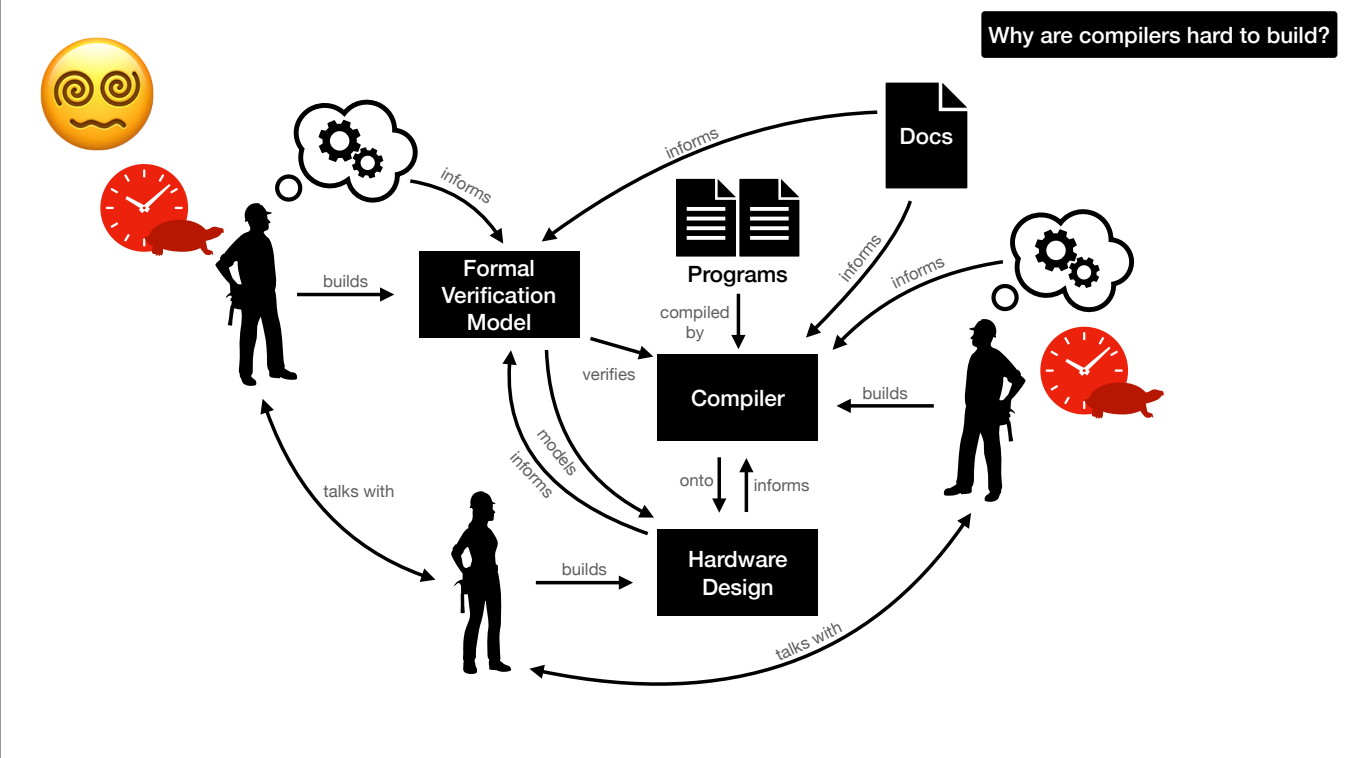
Why are compilers hard to build?

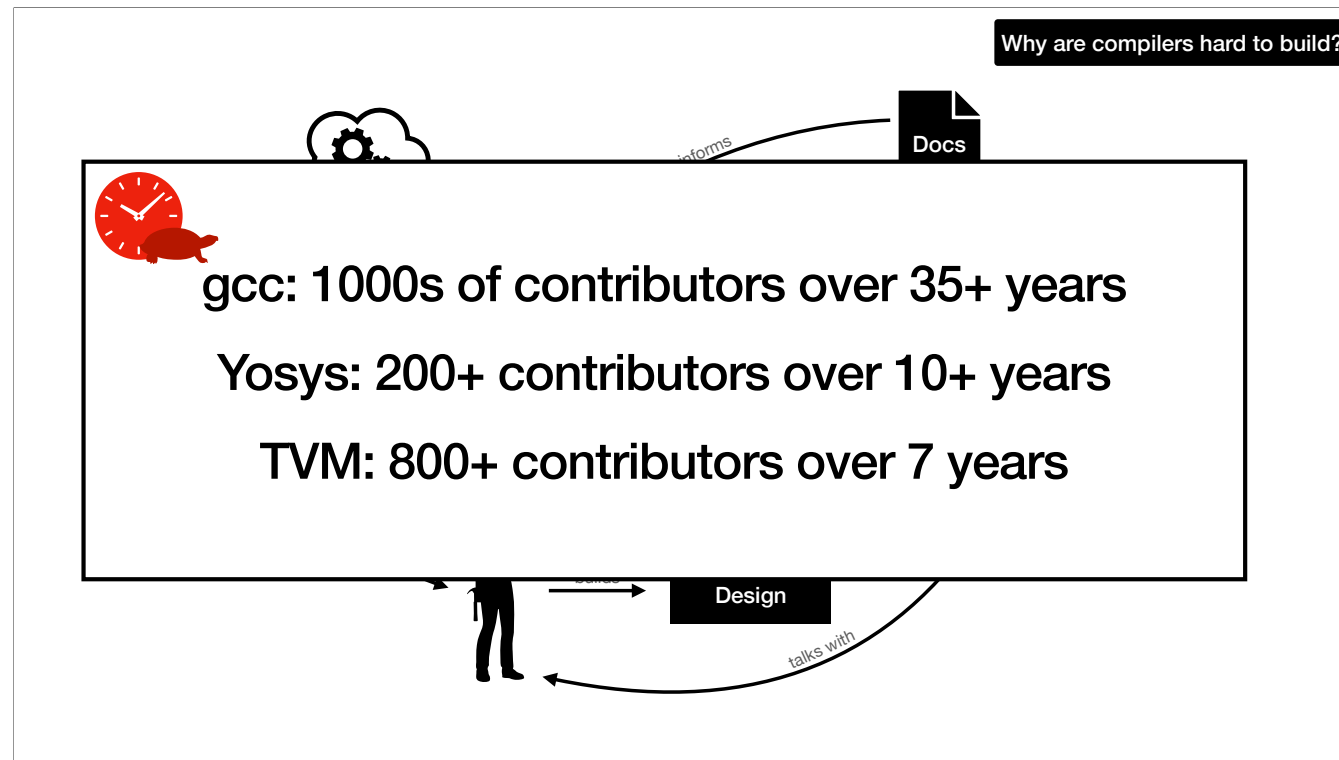


Why are compilers hard to build?

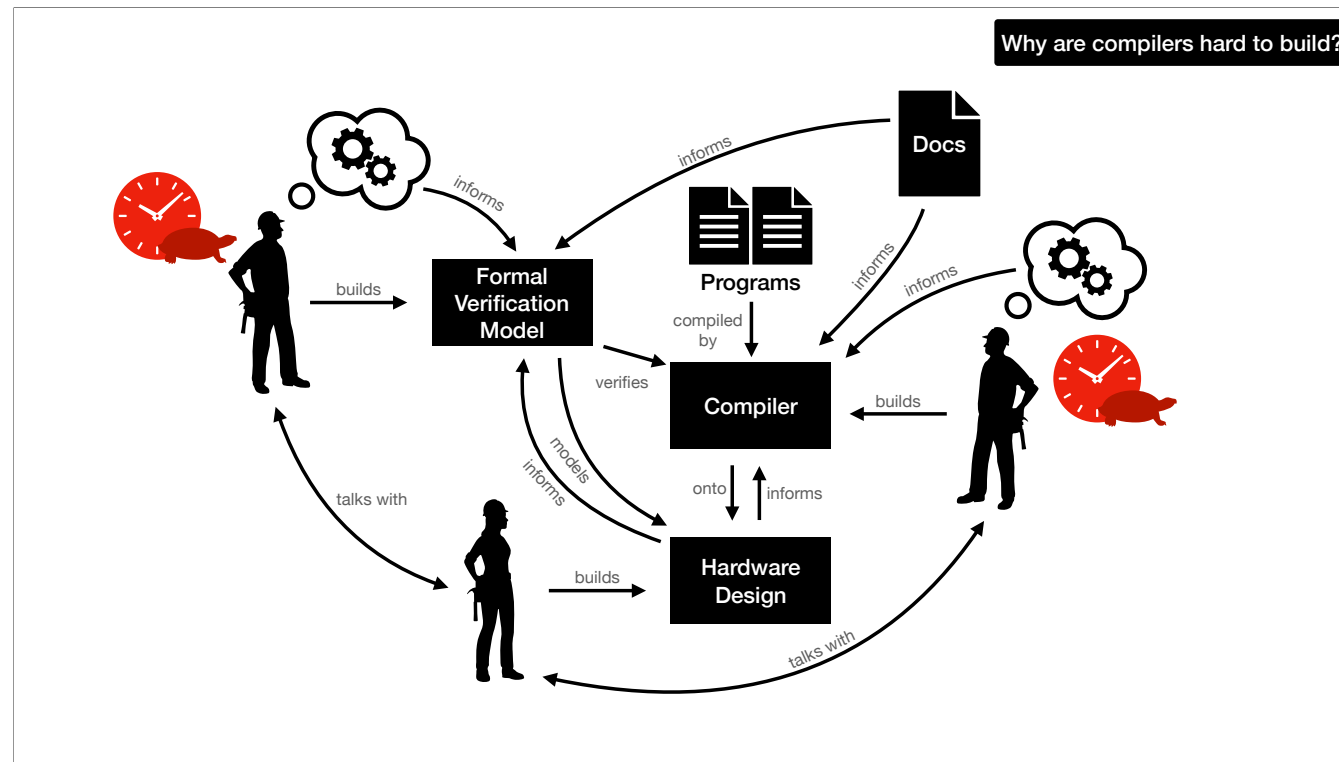









Which is clear once we see the sheer number of individual contributors and time that have gone into major open source compilers like gcc, the hardware synthesis tool Yosys, and the deep learning compiler TVM.



Furthermore, building a compiler is a [build] bug-prone process. In fact,



You can build a strong research career centered on finding and fixing bugs in these large, open-source compilers.

Why are compilers hard to build?



Finding and Understanding Bugs in C Compilers

Xuejun Yang Yang Chen Eric Eide John Regehr
University of Utah, School of Computing
{jxyang, chenyang, eeide, regehr}@cs.utah.edu
(Csmith)

Finds bugs
in gcc

Finding and Understanding Bugs in FPGA Synthesis Tools

Yann Herklotz John Wickerson
yann.herklotz15@imperial.ac.uk j.wickerson@imperial.ac.uk
Imperial College London Imperial College London
London, UK London, UK
(Verismith)

Finds bugs
in Yosys

A Comprehensive Study of Deep Learning Compiler Bugs

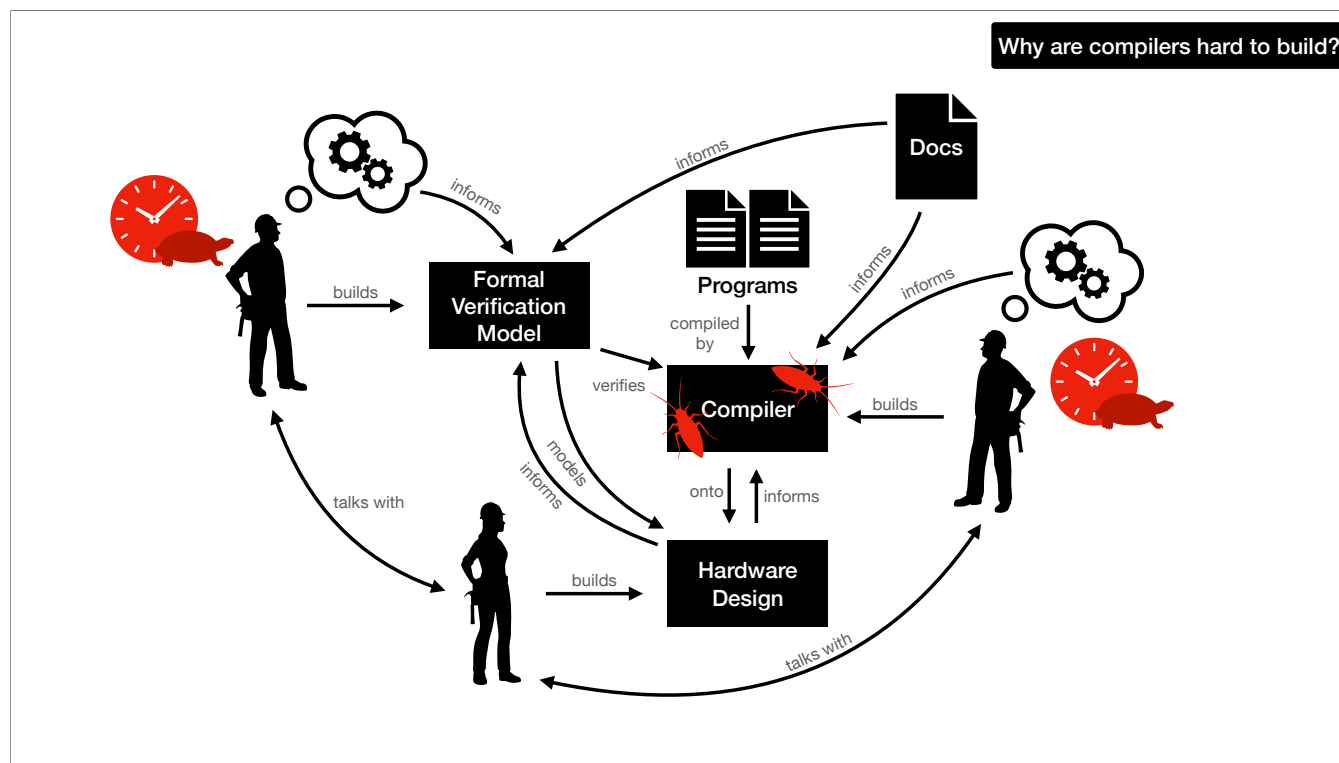
| | | |
|---|--|--|
| Qingchao Shen College of Intelligence and Computing, Tianjin University School of New Media and Communication, Tianjin University China qingchao@tju.edu.cn | Haoyang Ma College of Intelligence and Computing, Tianjin University China haoyang_9804@tju.edu.cn | Junjie Chen* College of Intelligence and Computing, Tianjin University China junjiechen@tju.edu.cn |
|---|--|--|

(et al.)

Finds bugs
in TVM




talks




And to make matters worse, these costs are multiplicative. That is,

Why are compilers hard to build?



...for each new piece of hardware,
[build] the entire process needs to be repeated to build a new compiler. Though there exists compiler frameworks such as LLVM and MLIR which lessen the burden on compiler engineers, the process still requires significant effort and expertise.

Why are compilers hard to build?



Why are compilers hard to build?

So we asked the question ...

What we've seen is that

[build] ...

And, importantly,

[build] ...

A natural question after all of this is

[build] ...

It might sound optimistic given how much effort it is to build a compiler, but let's at least entertain the possibility.

Why are compilers hard to build?

Building a compiler requires significant engineering effort and induces numerous bugs.

Why are compilers hard to build?

Building a compiler requires significant engineering effort and induces numerous bugs.


Those costs are multiplied with every new hardware design.

Why are compilers hard to build?

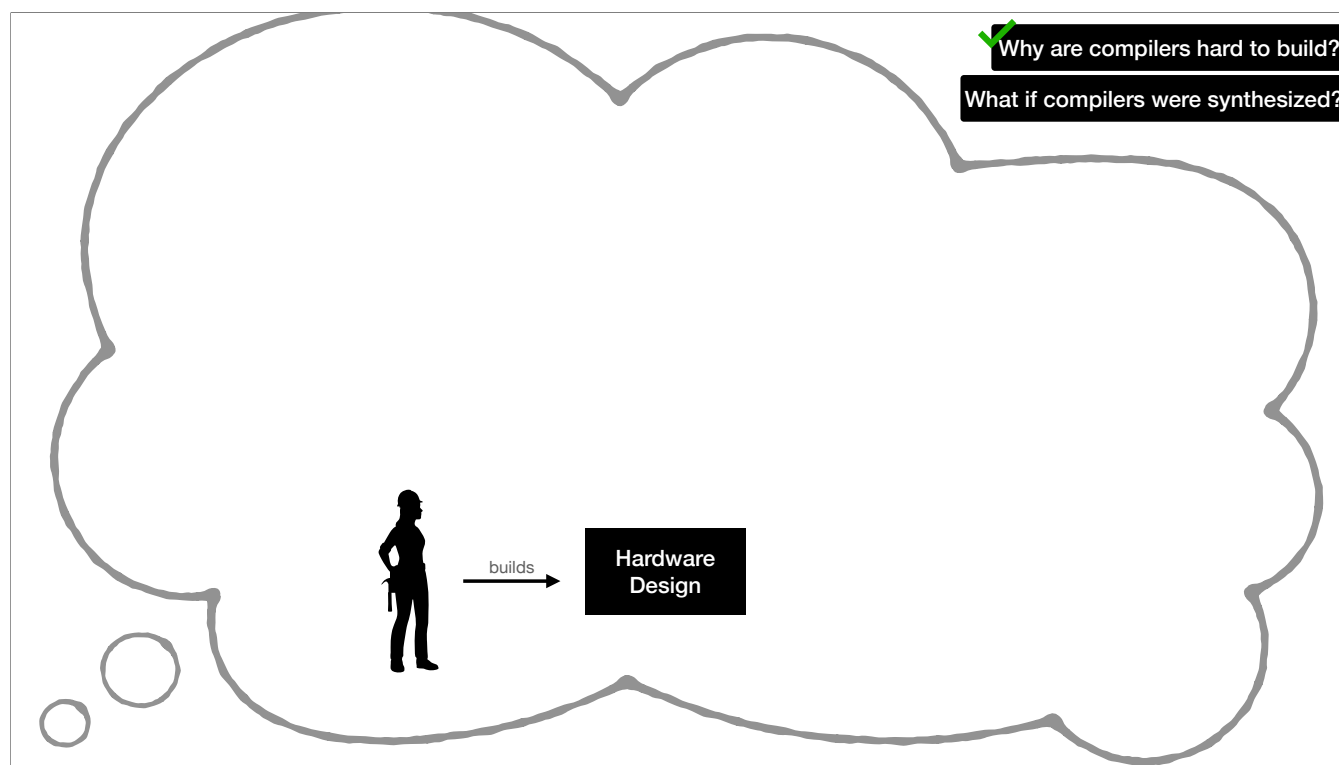
Building a compiler requires significant engineering effort and induces numerous bugs.

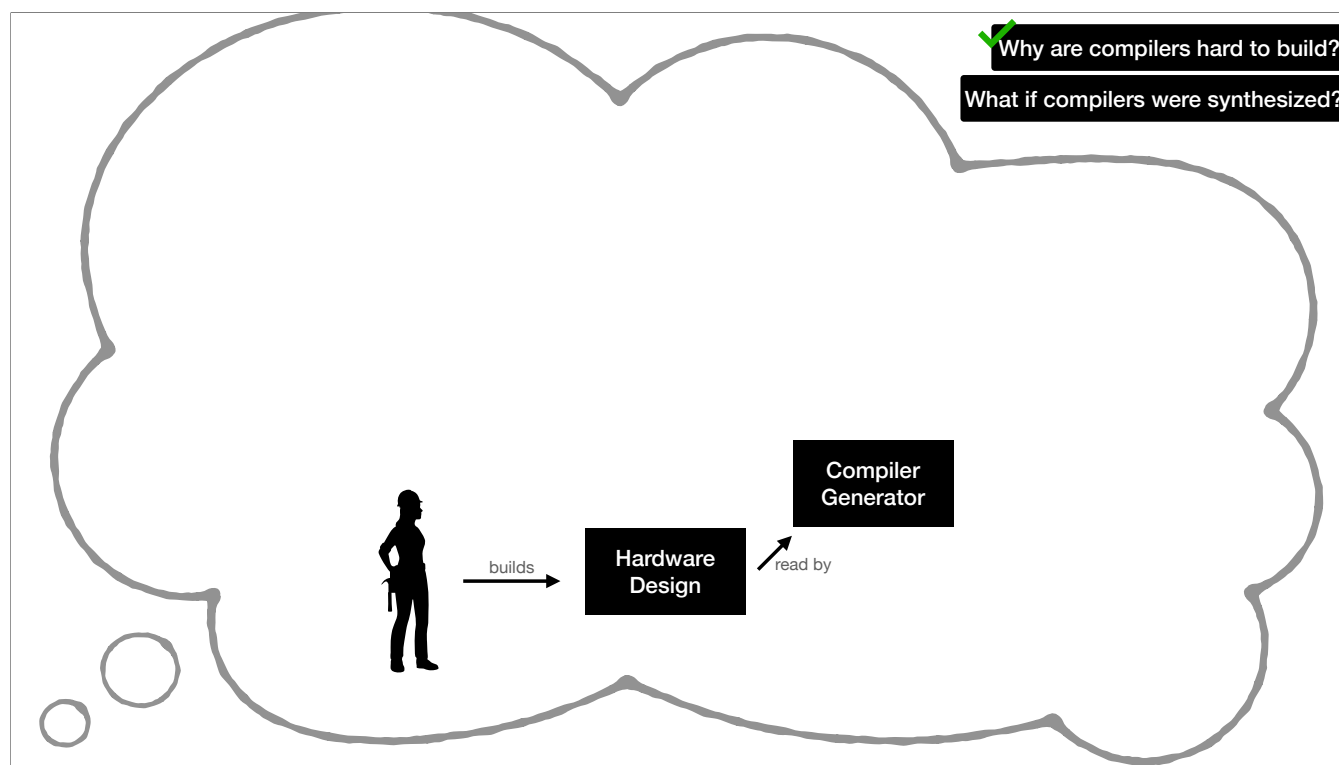
Those costs are multiplied with every new hardware design.

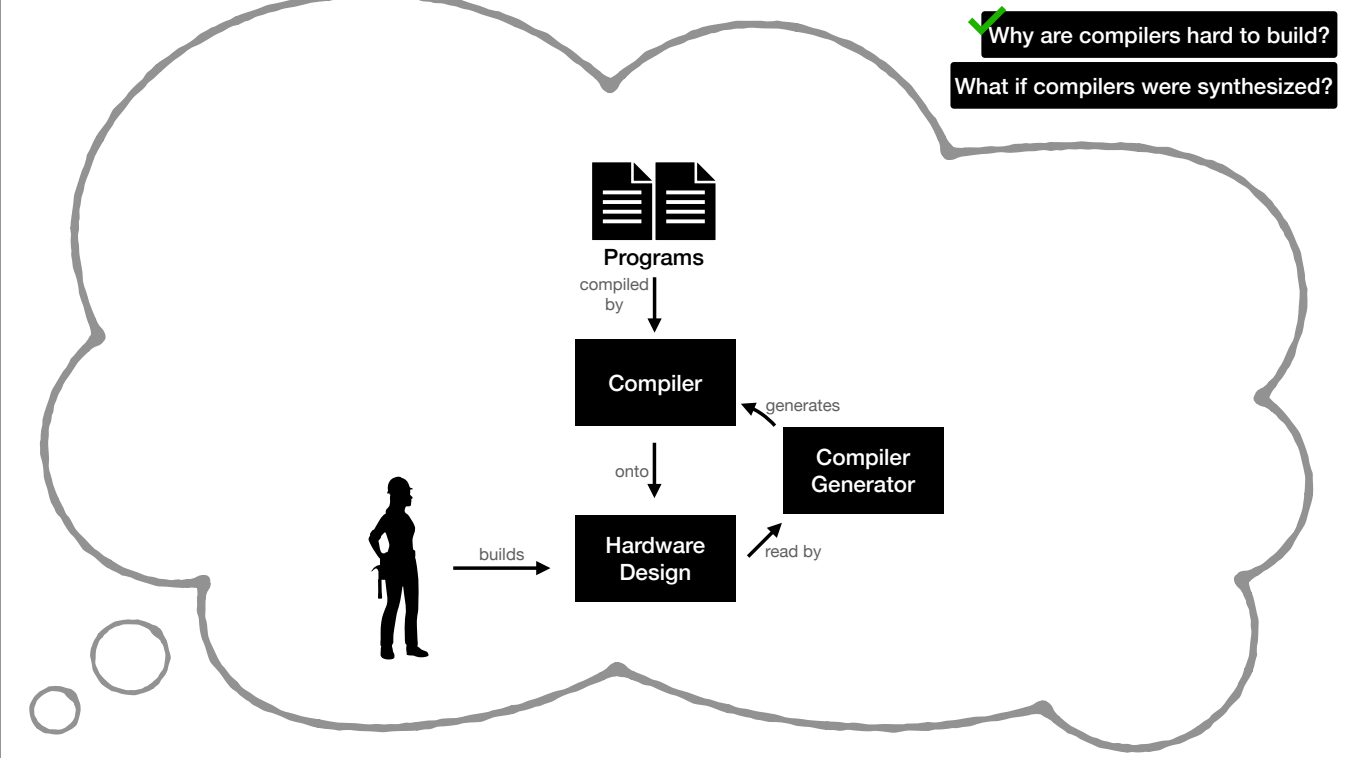
What if compilers were *synthesized*?
(i.e., automatically generated?)

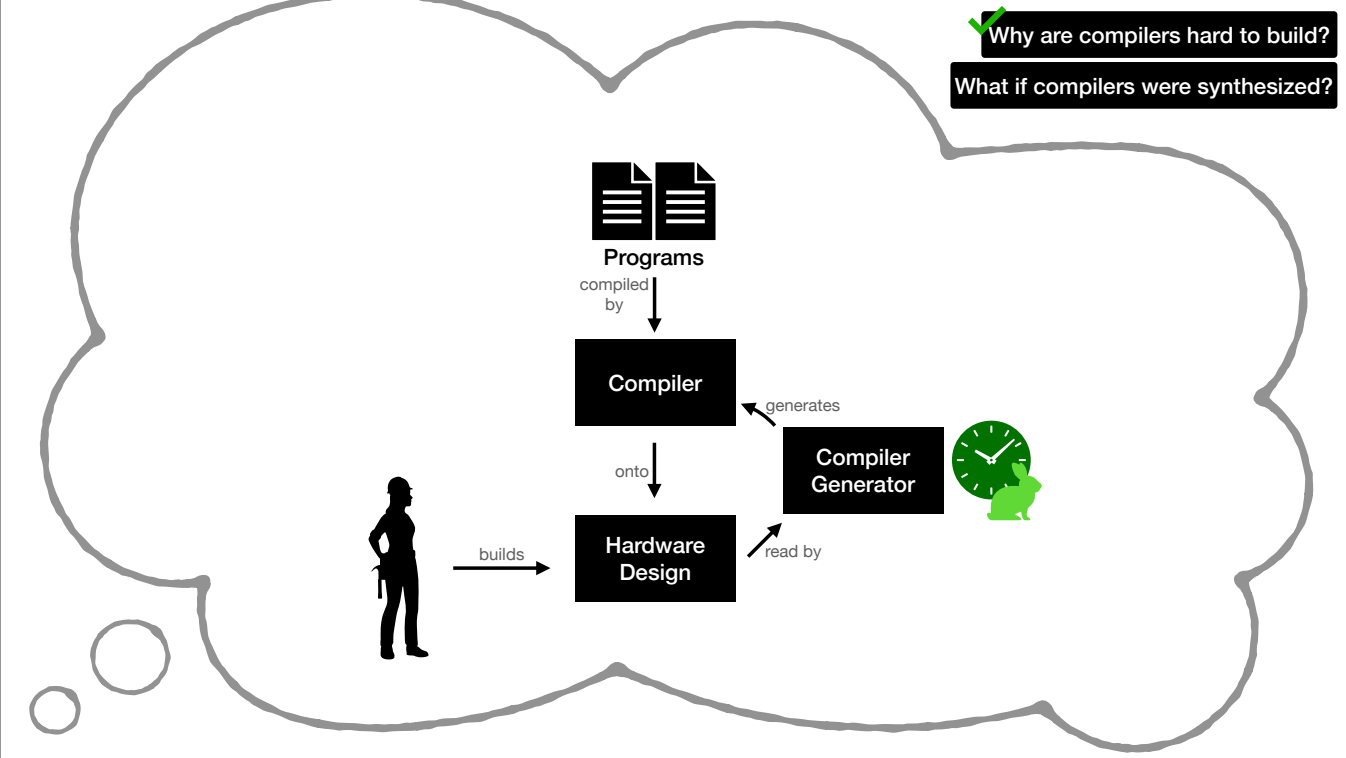


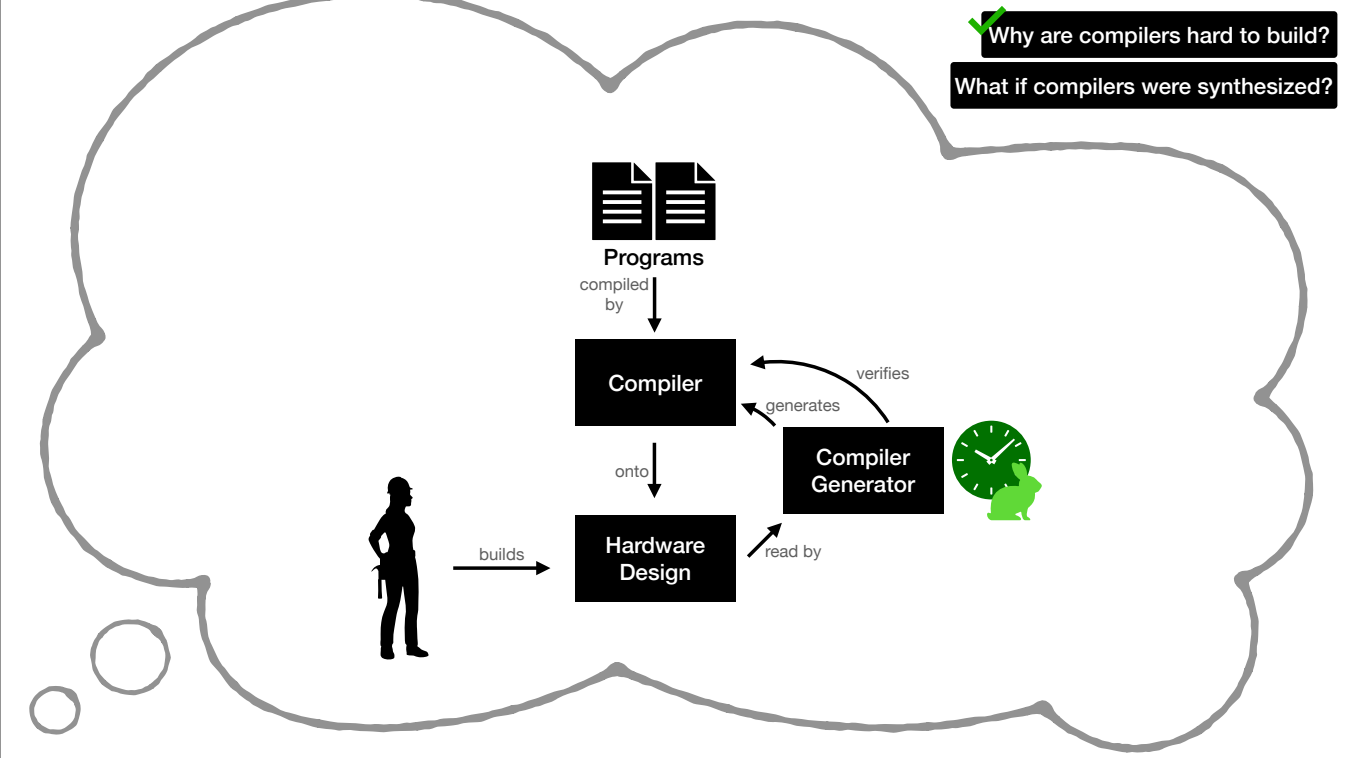
In the ideal case, our
[build] hardware designer still builds their hardware design. But now, that design is read by
[build] a compiler generator, which
[build] generates a compiler directly from their hardware implementation.
Automating compiler construction would save an immense amount of
[build] engineering effort and time.
In addition, depending on how the compiler generator is built, the generator could
[build] verify the compiler as it's being generated, producing
[build] a bug-free compiler, saving verification time and effort as well.

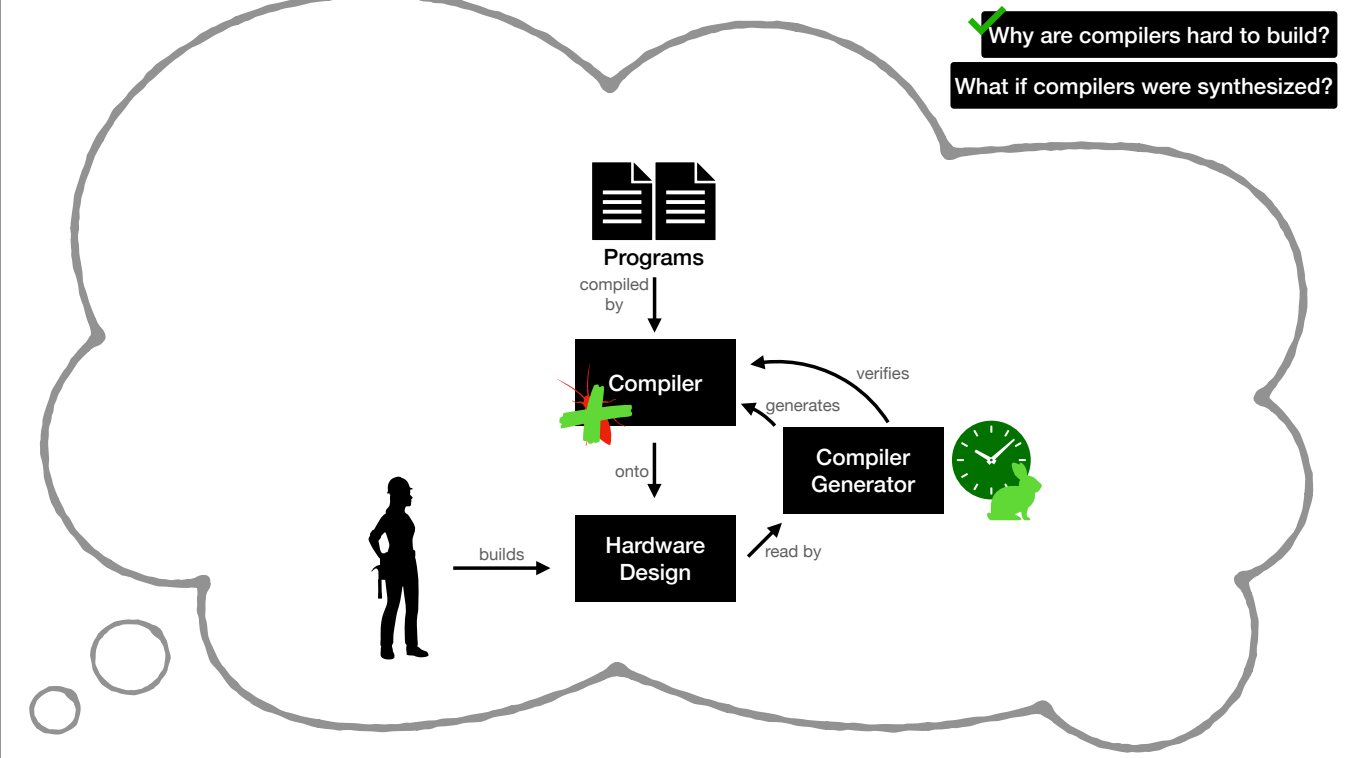







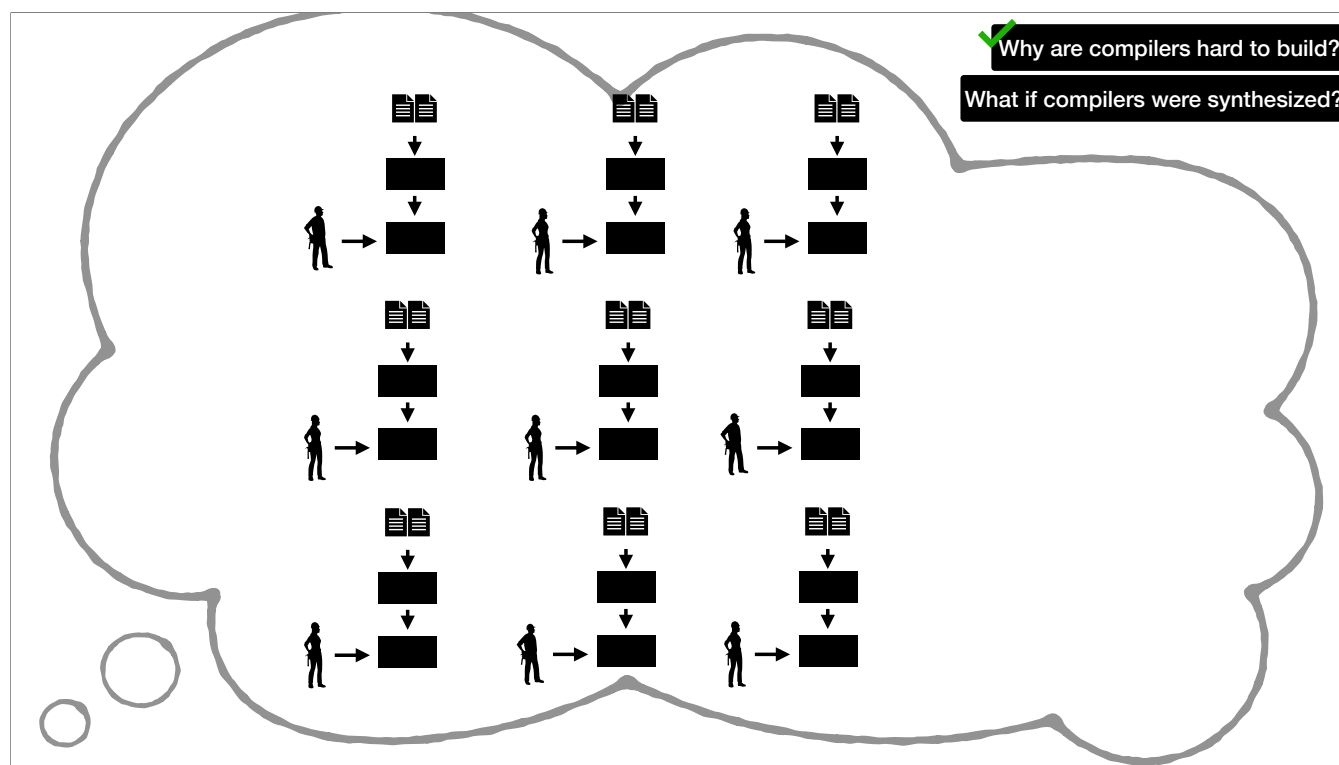


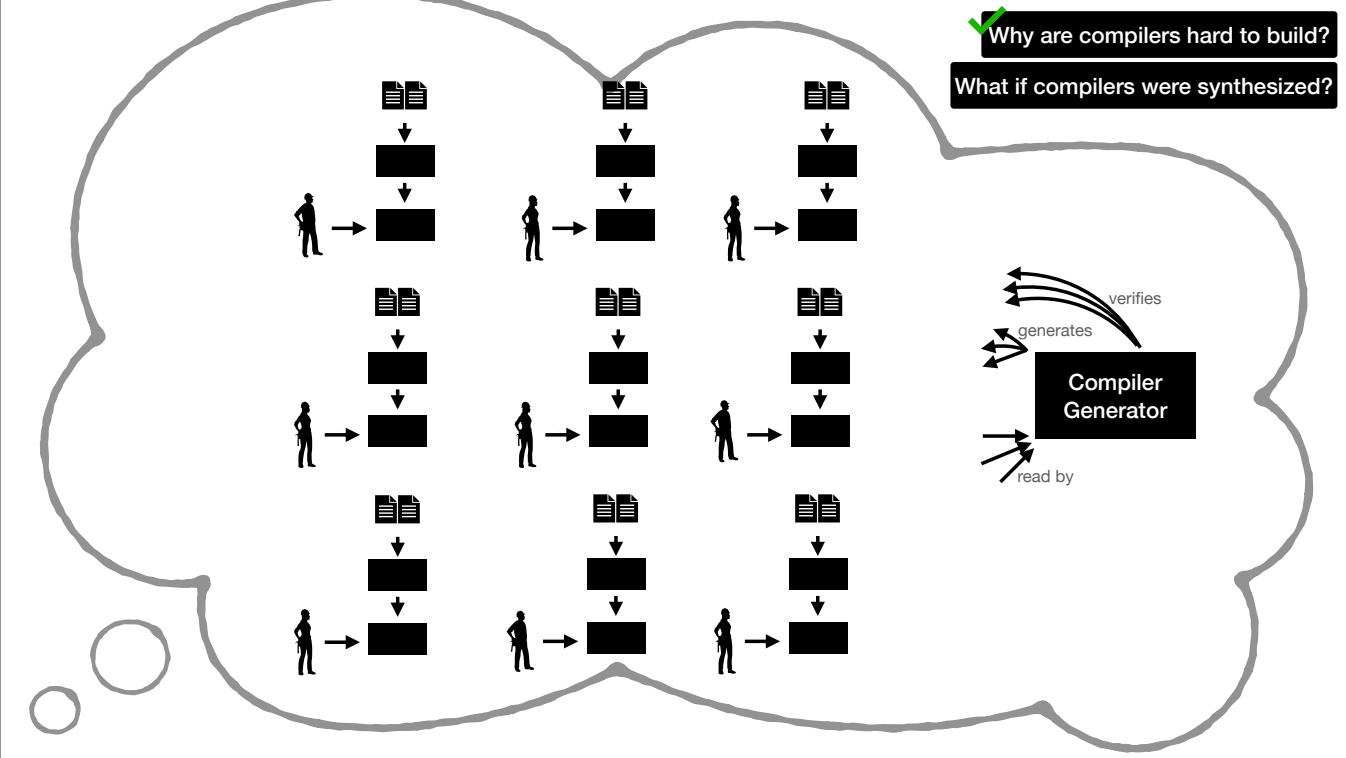







But best of all, this approach can scale
[build] as every new hardware design can reuse
[build] the same compiler generator, removing that multiplicative factor on effort and bugs that exists today.
In this ideal world,
[build] there is no hardware lottery, at least not because of compilers.







✓✓ Why are compilers hard to build?

What if compilers were synthesized?

So we see that
[build] ...
[build] ...

✓ Why are compilers hard to build?

What if compilers were synthesized?

Automatically generating compilers can reduce engineering effort and eliminate bugs.

What if compilers were synthesized?

Automatically generating compilers can reduce engineering effort and eliminate bugs.

Furthermore, the approach scales with new hardware designs, thus fighting against the hardware lottery!



With that, I will now introduce the thesis statement of this talk, in which I claim
[build] ...

**Compilers should be generated from formal
models of hardware.**

**Compilers should be generated from formal
models of hardware.**

**With the growing diversity of hardware and the
rapid improvement of automated reasoning,
now is the time to make this a reality.**

Generating Compilers → Why Now? → Case Study: Lakeroad → Call to Action

Here's our roadmap for the rest of the talk.

Generating Compilers → **Why Now?** → Case Study: Lakeroad → Call to Action

Let's talk about why now is the time to do this research.

Compilers should be generated from formal
models of hardware.

**With the growing diversity of hardware and the
rapid improvement of automated reasoning,
now is the time to make this a reality.**


We claim that, with the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make automatic generation of compilers a reality.

Compilers should be generated from formal models of hardware.


With the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make this a reality.

First, let's talk about what we mean by the growing diversity of hardware.


As soon as I mention the diversity of hardware, I'm sure that the first thing that pops into peoples' minds is hardware for machine learning, such as [build] GPUs and [build] custom machine learning ASICs. Yet even within [build] processors like Apple's A16, we're seeing the addition of specialized accelerators like GPUs and Neural Engines. Consider also platforms like [build] Xilinx's Zynq chip, which includes both an ARM CPU and a reconfigurable FPGA, making quite an interesting target for compilers. Lastly, far from the realm of silicon-based computing, people have begun computing using things like [build] DNA strand displacement or [build] metamaterials. Though these are far from what we would normally consider "hardware", they require compilers nonetheless! Given the dizzying array of hardware available today, it's clear that [build] hardware is growing more diverse, and that compilers for new hardware are desperately needed. This diversity can be intimidating: [build] how could we possibly generate compilers for all of this hardware? But the explosion of new hardware platforms actually works in our favor, [build] because as hardware diversifies, it gets more specialized, and thus, potentially easier to target with automated methods.




NVIDIA Tensor Cores




NVIDIA Tensor Cores




Google TPU




AWS Inferentia



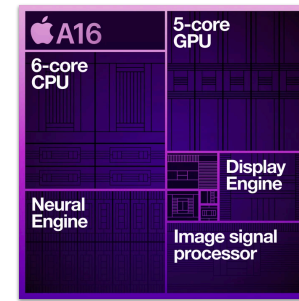
NVIDIA Tensor Cores




Google TPU




AWS Inferentia




Apple A16




NVIDIA Tensor Cores



Google TPU




AWS Inferentia




Apple A16




Xilinx Zynq




NVIDIA Tensor Cores



Google TPU




AWS Inferentia




Apple A16




Xilinx Zynq




Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.



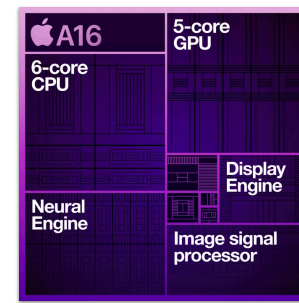
NVIDIA Tensor Cores



Google TPU




AWS Inferentia




Apple A16




Xilinx Zynq




Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.




Parsa et al. Universal Mechanical Polycomputation in Granular Matter.




NVIDIA Tensor




Google TPU




Hardware is growing more diverse; more compilers are desperately needed!




Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.




Parsa et al. Universal Mechanical Polycomputation in Granular Matter.



NVIDIA Tensor




Google TPU




Hardware is growing more diverse; more compilers are desperately needed!


How could we possibly support all of this hardware?




Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.




Parsa et al. Universal Mechanical Polycomputation in Granular Matter.



NVIDIA Tensor




Google TPU




Hardware is growing more diverse; more compilers are desperately needed!

How could we possibly support all of this hardware?


As as hardware diversifies, it gets more specialized, and thus easier to target!



Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.




Parsa et al. Universal Mechanical Polycomputation in Granular Matter.



Previous work on automatically generating compilers largely focuses on
[build] processors.

Generating compilers for processors is

[build] quite a difficult task, as processors are general purpose, and compilers for general purpose processors must handle all of their capabilities.



NVIDIA Tensor Cores

Describing Instruction Set Processors Using nML

A. Fauth¹ J. Van Praet² M. Freericks¹

¹Institut für Technische Informatik
Tech. Univ. Berlin, Franklinstr. 28/29
D-10587 Berlin, Germany

²IMEC
Kapeldreef 75
B-3001 Leuven, Belgium

1995

Retargetable Generation of Code Selectors from HDL Processor Models


University of Dortmund, Dept. of Computer Science 12, 44221 Dortmund, Germany
email: leupers|marwedel@ls12.informatik.uni-dortmund.de

1997

Automatic Tool Generation from Structural Processor Descriptions

Florian Brandner


2009



Gerasimova et al. *Connect*

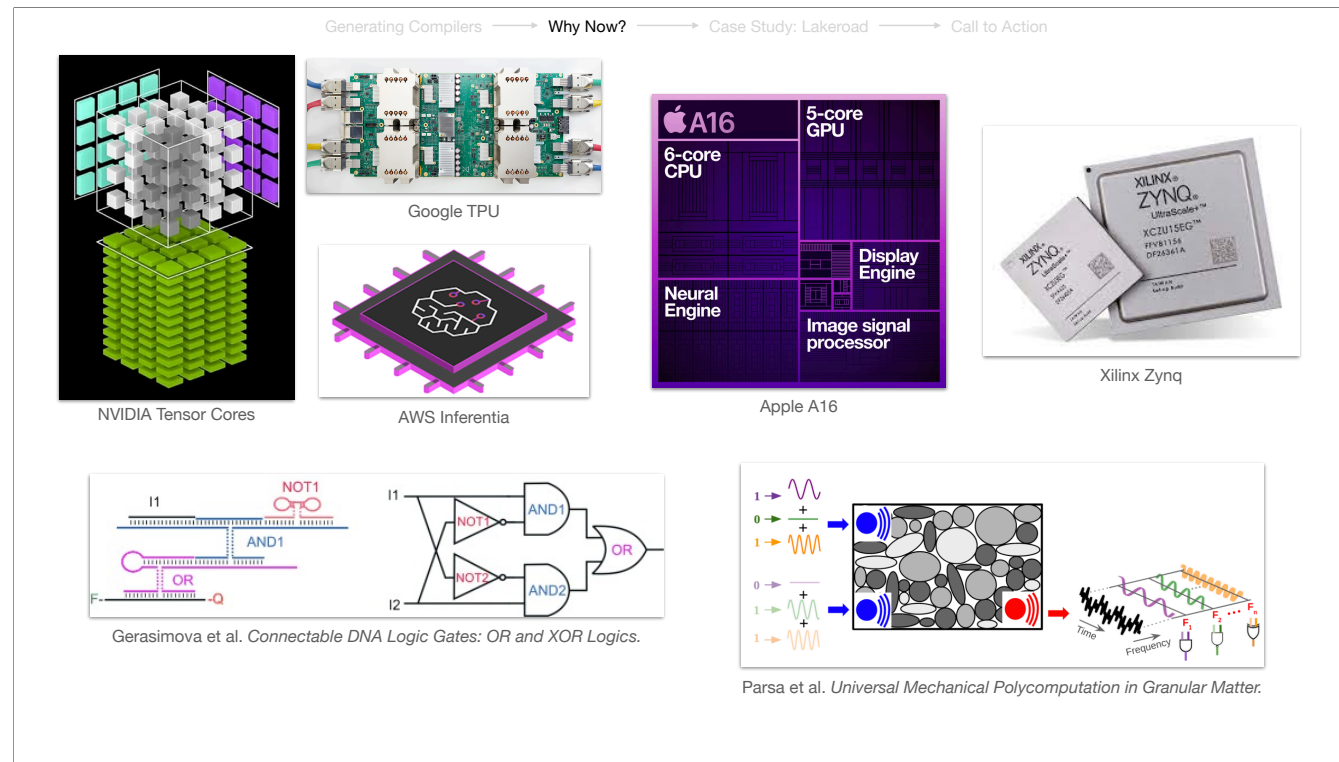


Xilinx Zynq



Parsa et al. *Universal Mechanical Polycomputation in Granular Matter*.


Generating compilers for general-purpose hardware is difficult.




On the other hand, all of the new hardware we're talking about is special purpose,

[build] which makes the task of reasoning about hardware's behavior much more feasible for automated methods.

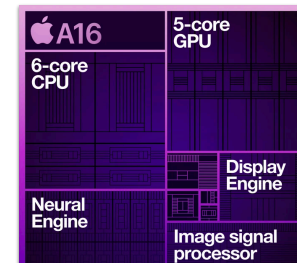
But not only is our hardware more amenable to automated reasoning; our tools for automated reasoning are now powerful enough to take on the task of automated compiler generation.



NVIDIA Tensor




Google TPU




ynq

Specialized hardware is easier to target with automated reasoning tools!



Gerasimova et al. Connectable DNA Logic Gates: OR and XOR Logics.



Parsa et al. Universal Mechanical Polycomputation in Granular Matter.


Consider, for example,
[build] SAT and SMT solvers, whose performance have been
[build] steadily increasing,

Or the relatively new technique of
[build] equality saturation, which has already shown great promise for compiler construction,

And of course it wouldn't be a talk in 2023 if I didn't mention
[build] large language models, which are powerful tools for generating hardware code, among other tasks.


Given that this is only a small selection of the automated reasoning tools available, it's clear that
[build] ...

SAT/SMT




Järvisalo et al. 2012. The international SAT solver competitions.

SAT/SMT



Järvisalo et al. 2012. The international SAT solver competitions.

SAT/SMT



Järvisalo et al. 2012. The international SAT solver competitions.

Equality Saturation

Vectorization for Digital Signal Processors
via Equality Saturation

Alexa VanHattum
Cornell University
Ithaca, NY, USA


Rachit Nigam
Cornell University
Ithaca, NY, USA

Vincent T. Lee
Facebook Reality Labs Research
Redmond, WA, USA

James Bornholt
The University of Texas at Austin
Austin, TX, USA

Adrian Sampson
Cornell University
Ithaca, NY, USA

SAT/SMT



Equality Saturation

Vectorization for Digital Signal Processors
via Equality Saturation

Alexa VanHattum
Cornell University
Ithaca, NY, USA

Rachit Nigam
Cornell University
Ithaca, NY, USA

Vincent T. Lee
Facebook Reality Labs Research
Redmond, WA, USA

James Bornholt
The University of Texas at Austin
Austin, TX, USA


Adrian Sampson
Cornell University
Ithaca, NY, USA

Benchmarking Large Language Models for
Automated Verilog RTL Code Generation

Shailja Thakur*, Baleegh Ahmad*, Zhenxing Fan*, Hammond Pearce*,
Benjamin Tan[†], Ramesh Karri*, Brendan Dolan-Gavitt*, Siddharth Garg*
*New York University, [†]University of Calgary

ML/LLMs

SAT/SMT



Järvisalo et al. 2012. The international SAT solver competitions.

Equality Saturation

Vectorization for Digital Signal Processors
via Equality Saturation

Alexa VanHattum
Cornell University
Ithaca, NY, USA

Rachit Nigam
Cornell University
Ithaca, NY, USA

Vincent T. Lee
Facebook Reality Labs Research
Redmond, WA, USA

James Bornholt
The University of Texas at Austin

Adrian Sampson
Cornell University

Automated reasoning tools are ready for the task of compiler generation.

Models for
Automated Verilog RTL Code Generation

Shailja Thakur*, Baleegh Ahmad*, Zhenxing Fan*, Hammond Pearce*,
Benjamin Tan[†], Ramesh Karri*, Brendan Dolan-Gavitt*, Siddharth Garg*
*New York University, [†]University of Calgary

ML/LLMs

Compilers should be generated from formal
models of hardware.

**With the growing diversity of hardware and the
rapid improvement of automated reasoning,
now is the time to make this a reality.**

So, why is now the time to make the automatic generation of compilers a reality? Well,

[build] ...

Furthermore,

[build] ...

Finally,

[build] ...

Compilers should be generated from formal models of hardware.

With the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make this a reality.

Hardware is diversifying, and we need new compilers.

Compilers should be generated from formal
models of hardware.

**With the growing diversity of hardware and the
rapid improvement of automated reasoning,
now is the time to make this a reality.**

Hardware is diversifying, and we need new compilers.
Modern targets are more amenable to automated methods.

Compilers should be generated from formal
models of hardware.

**With the growing diversity of hardware and the
rapid improvement of automated reasoning,
now is the time to make this a reality.**

Hardware is diversifying, and we need new compilers.
Modern targets are more amenable to automated methods.
Automated reasoning tools are ready for the task of compiler generation.

Generating Compilers → **Why Now?** → Case Study: Lakeroad → Call to Action

Now,

Generating Compilers → Why Now? → **Case Study: Lakeroad** → Call to Action

...let's talk about a concrete example of generating a compiler from hardware models, in a project we call Lakeroad.

Compilers should be generated from formal models of hardware.


With the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make this a reality.

We keep talking about the growing diversity of hardware platforms. Now, [build] let's look at a concrete example: FPGAs.

Compilers should be generated from formal models of hardware.

With the growing diversity of hardware and the rapid increase in automated reasoning, now is the time to make this a reality.

Let's look at a concrete example: FPGAs.



FPGAs are reconfigurable devices that can be used to implement hardware.


At a high level, though, you can think of an FPGA as being a bag filled with parts, or primitives.

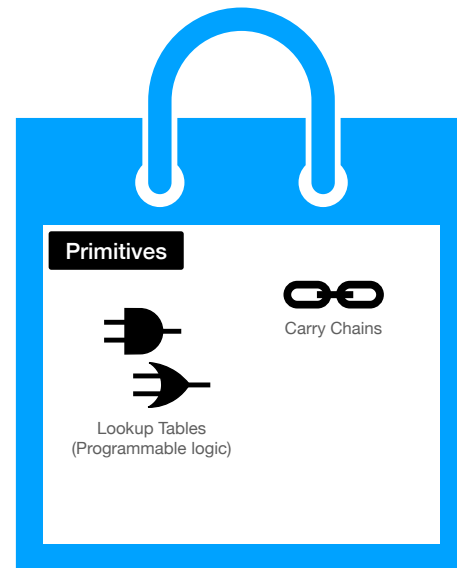


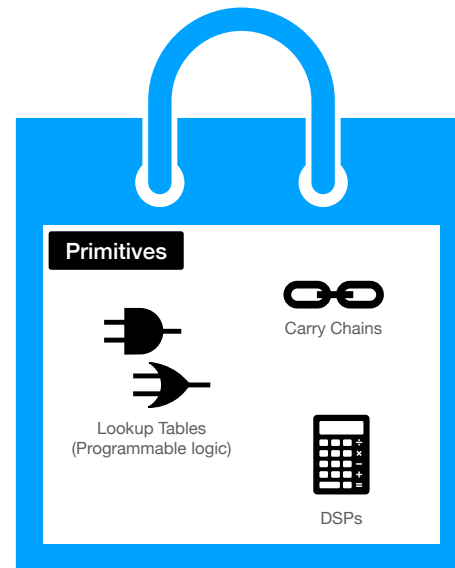
In the past, FPGAs consisted only of
[build] lookup tables, which are primitives that can be configured to implement logic gates.

Over the years, FPGAs have added specialized primitives such as
[build] carry chains to implement fast arithmetic.

One of the most interesting and impactful additions to FPGAs has been the inclusion of
[build] digital signal processors or DSPs, which are small, programmable embedded processors.







Compilers should be generated from formal models of hardware.

With the growing diversity of hardware and the rapid growth of formalized reasoning, this is a reality.

Even *within* FPGAs, hardware is diversifying.

So we can see that

[build] ...

[build] Are new primitives a challenge for FPGA compilers, as our thesis would suggest?

Compilers should be generated from formal models of hardware.

With the growing diversity of hardware and the rapid pace of formalized reasoning, this is a reality.

Even *within* FPGAs, hardware is diversifying.
Are new primitives a challenge for FPGA compilers?

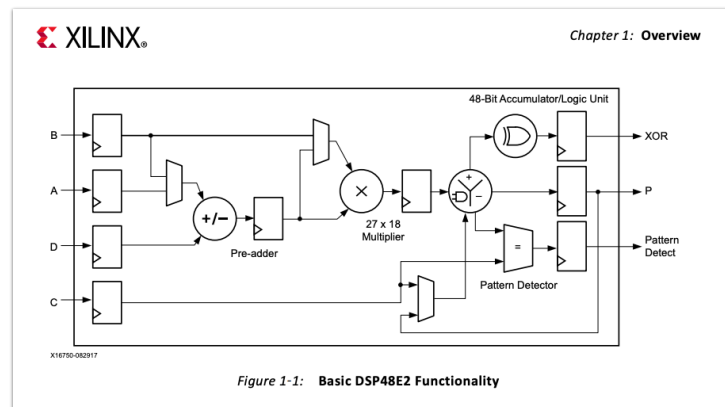
Are new primitives a challenge for FPGA compilers?

$$((d + a) * b) ^ c$$

To test this, we will attempt to compile a simple hardware design onto a Xilinx FPGA's DSP.

In our case, our simple design takes four inputs and computes this expression in three pipeline stages.

Are new primitives a challenge for FPGA compilers?



$$((d + a) * b) \wedge c$$

The Xilinx DSP documentation claims that this expression is supported on the DSP. But when we attempt to compile our design using the state of the art compiler for Xilinx FPGAs, we see a surprising result.

Are new primitives a challenge for FPGA compilers?

Report Cell Usage:

| | Cell | Count |
|---|---------|-------|
| 1 | DSP48E1 | 2 |
| 2 | LUT2 | 10 |
| 3 | SRL16E | 10 |
| 4 | FDRE | 10 |

We see that the design is in fact compiled onto DSPs; but instead of [build] using a single DSP as expected, it uses two DSPs and ten look up tables.

Are new primitives a challenge for FPGA compilers?

Report Cell Usage:

| | Cell | Count |
|---|---------|-------|
| 1 | DSP48E1 | 2 |
| 2 | LUT2 | 10 |
| 3 | SRL16E | 10 |
| 4 | FDRE | 10 |

Should use only
a single DSP!

Are new primitives a challenge for FPGA compilers?

So, are new primitives...

From our brief experiment, it seems like the answer is

[build] yes!

[build] But this is unsurprising...

Are new primitives a challenge for FPGA compilers?

On brief inspection, yes!

Are new primitives a challenge for FPGA compilers?

On brief inspection, yes!

But this is unsurprising—DSPs are complicated.

Generating Compilers → Why Now? → **Case Study: Lakeroad** → Call to Action

The manual for our Xilinx DSP alone is
[build] over 75 pages long.



| | |
|---|----|
| MULTISIGNOUT and CARRYCASCOUT | 72 |
| Summary | 73 |
| Appendix A: Additional Resources and Legal Notices | |
| Xilinx Resources | 75 |
| Solution Centers | 75 |
| Documentation Navigator and Design Hubs | 75 |
| References | 76 |
| Please Read: Important Legal Notices | 77 |



| | |
|---|----|
| MULTISIGNOUT and CARRYCASCOUT | 72 |
| Summary | 73 |
| Appendix A: Additional Resources and Legal Notices | |
| Xilinx Resources | 75 |
| Solution Centers | 75 |
| Documentation Navigator and Design Hub | 75 |
| References | 76 |
| Please Read: Important Legal Notices | 77 |

DSP manual is over 75
pages long

```
module DSP48E2 #(
    parameter integer ACASCREG = 1,
    parameter integer ADREG = 1,
    parameter integer ALUMODEREG = 1,
    parameter AMULTSEL = "A",
    parameter integer AREG = 1,
    parameter AUTORESET_PATDET = "NO_RESET",
    parameter AUTORESET_PRIORITY = "RESET",
    parameter A_INPUT = "DIRECT",
    ...
)(
    output [29:0] ACOUT,
    output [17:0] BCOUT,
    output CARRYCASCOUT,
    ...

    input [29:0] A,
    input [29:0] ACIN,
    input [3:0] ALUMODE,
    input [17:0] B,
    input [17:0] BCIN,
    input [47:0] C,
    ...
);
```

Furthermore, instantiating a DSP requires

[build] setting over 100 ports and parameters, while

[build] obeying strict requirements on port and parameter values, described throughout the 75 page manual.

With all of these complex interdependencies on what values are legal for which parameters based on the values of other parameters,

[build] configuring a DSP is starting to sound a lot like writing a program.

```
module DSP48E2 #(
  parameter integer ACASCREG = 1,
  parameter integer ADREG = 1,
  parameter integer ALUMODEREG = 1,
  parameter AMULTSEL = "A",
  parameter integer AREG = 1,
  parameter AUTORESET_PATDET = "NO_RESET",
  parameter AUTORESET_PRIORITY = "RESET",
  parameter A_INPUT = "DIRECT",
  ...
)(
  output [29:0] ACOUT,
  output [17:0] BCOUT,
  output CARRYCASCOUT,
  ...

  input [29:0] A,
  input [29:0] ACIN,
  input [3:0] ALUMODE,
  input [17:0] B,
  input [17:0] BCIN,
  input [47:0] C,
  ...
);
```

Configuring the DSP
requires setting 100+
ports and parameters

```
module DSP48E2 #(
    parameter integer ACASCREG = 1,
    parameter integer ADREG = 1,
    parameter integer ALUMODEREG = 1,
    parameter AMULTSEL = "A",
    parameter integer AREG = 1,
    parameter AUTORESET_PATDET = "NO_RESET",
    parameter AUTORESET_PRIORITY = "RESET",
    parameter A_INPUT = "DIRECT",
    ...
)(
    output [29:0] ACOUT,
    output [17:0] BCOUT,
    output CARRYCASCOUT,
    ...

    input [29:0] A,
    input [29:0] ACIN,
    input [3:0] ALUMODE,
    input [17:0] B,
    input [17:0] BCIN,
    input [47:0] C,
    ...
);
```

Configuring the DSP
requires setting 100+
ports and parameters

Table 2-4: OPMODE Control Bits Select X Multiplexer Outputs

| W OPMODE[8:7] | Z OPMODE[6:4] | Y OPMODE[3:2] | X OPMODE[1:0] | X Multiplexer Output | Notes |
|------------------|------------------|------------------|------------------|-------------------------|--------------------------------------|
| xx | xxx | xx | 00 | 0 | Default |
| xx | xxx | 01 | 01 | M | Must select with OPMODE[3:2] = 01 |
| xx | xxx | xx | 10 | P | Requires PREG = 1 |
| xx | xxx | xx | 11 | A:B | 48-bits wide |

When either TWO24 or FOUR12 mode is selected, the
multiplier must not be used, and USE_MULT must be set
to NONE.

Notes:

- When these data pins are not used and to reduce leakage power dissipation, the data pin input signals must be tied High, the input register must be selected, and the CE and RST input control signals must be tied Low. An example of unused C input recommended settings would be setting C[47:0] = all ones, CREG = 1, CEC = 0, and RSTC = 0.
- These signals are dedicated routing paths internal to the DSP48E2 column. They are not accessible via general routing resources.
- All signals are active High.

```

module DSP48E2 #(
    parameter integer ACASCREG = 1,
    parameter integer ADREG = 1,
    parameter integer ALUMODEREG = 1,
    parameter AMULTSEL = "A",
    parameter integer AREG = 1,
    parameter AUTORESET_PATDET = "NO_RESET",
    parameter AUTORESET_PRIORITY = "RESET",
    parameter A_INPUT = "DIRECT",
    ...
)(
    output [29:0] ACOUT,
    output [17:0] BCOUT,
    output CARRYCASCOUT,
    ...

    input [29:0] A,
    input [29:0] ACIN,
    input [3:0] ALUMODE,
    input [17:0] B,
    input [17:0] BCIN,
    input [47:0] C,
    ...
);

```

Configuring the DSP
requires setting 100+
ports and parameters

Table 2-4: OPMODE Control Bits Select X Multiplexer Outputs

| W OPMODE[8:7] | Z OPMODE[6:4] | Y OPMODE[3:2] | X OPMODE[1:0] | X Multiplexer Output | Notes |
|------------------|------------------|------------------|------------------|-------------------------|--------------------------------------|
| xx | xxx | xx | 00 | 0 | Default |
| xx | xxx | 01 | 01 | M | Must select with OPMODE[3:2] = 01 |
| xx | xxx | xx | 10 | P | Requires PREG = 1 |
| xx | xxx | xx | 11 | A:B | 48-bits wide |

When either TWO24 or FOUR12 mode is selected, the
multiplier must not be used, and USE_MULT must be set
to NONE.

Notes:

- When these data pins are not used and to reduce leakage power dissipation, the data pin input signals must be tied High, the input register must be selected, and the CE and RST input control signals must be tied Low. An example of unused C input recommended settings would be setting C[47:0] = all ones, CREG = 1, CEC = 0, and RSTC = 0.
- These signals are dedicated routing paths internal to the DSP48E2 column. They are not accessible via general routing resources.
- All signals are active High.

Configuring a DSP sounds a lot like writing a program!

[build] so why not use program synthesis?

For those unfamiliar with program synthesis, I won't go into too much detail, but just know that [build] *solver aided program synthesis* is the process of...

```

module DSP48E2 #(
  parameter integer ACASCREG = 1,
  parameter integer ADREG = 1,
  parameter integer ALUMODEREG = 1,
  ...
  )(
  out
  out
  out
  ...
  inp
  inp
  inp
  inp
  inp
  input [4:0] C,
  ...
  );
    
```

Table 2-4: OPMODE Control Bits Select X Multiplexer Outputs

| W | Z | Y | X | X Multiplexer Output | Notes |
|-------------|-------------|-------------|-------------|----------------------|-------|
| OPMODE[8:7] | OPMODE[6:4] | OPMODE[3:2] | OPMODE[1:0] | | |

Insight #1: configuring DSPs and other complex primitives is similar to writing a program...

...so use *program synthesis*.

High
ed C
g


```
module DSP48E2 #(
  parameter integer ACASCREG = 1,
  parameter integer ADREG = 1,
  parameter integer ALUMODEREG = 1,
  ...
  out
  out
  out
  ...
  inp
  inp
  inp
  inp
  input [7:0] C,
  ...
);
```

Table 2-4: OPMODE Control Bits Select X Multiplexer Outputs

| W | Z | Y | X | X Multiplexer Output | Notes |
|-------------|-------------|-------------|-------------|----------------------|-------|
| OPMODE[8:7] | OPMODE[6:4] | OPMODE[3:2] | OPMODE[1:0] | | |

Insight #1: configuring DSPs and other complex primitives is similar to writing a program...

...so use *program synthesis*.

Solver-aided program synthesis: using SMT/SAT/etc. to generate programs by solving a set of constraints.

```

module DSP48E2 #(
    parameter integer ACASCREG = 1,
    parameter integer ADREG = 1,
    parameter integer ALUMODEREG = 1,
    parameter AMULTSEL = "A",
    parameter integer AREG = 1,
    parameter AUTORESET_PATDET = "NO_RESET",
    parameter AUTORESET_PRIORITY = "RESET",
    parameter A_INPUT = "DIRECT",
    ...
)(
    output [29:0] ACOUT,
    output [17:0] BCOUT,
    output CARRYCASCOUT,
    ...

    input [29:0] A,
    input [29:0] ACIN,
    input [3:0] ALUMODE,
    input [17:0] B,
    input [17:0] BCIN,
    input [47:0] C,
    ...
);

```

Table 2-4: OPMODE Control Bits Select X Multiplexer Outputs

| W OPMODE[8:7] | Z OPMODE[6:4] | Y OPMODE[3:2] | X OPMODE[1:0] | X Multiplexer Output | Notes |
|------------------|------------------|------------------|------------------|-------------------------|--------------------------------------|
| xx | xxx | xx | 00 | 0 | Default |
| xx | xxx | 01 | 01 | M | Must select with OPMODE[3:2] = 01 |
| xx | xxx | xx | 10 | P | Requires PREG = 1 |
| xx | xxx | xx | 11 | A:B | 48-bits wide |

When either TWO24 or FOUR12 mode is selected, the multiplier must not be used, and USE_MULT must be set to NONE.

Notes:

- When these data pins are not used and to reduce leakage power dissipation, the data pin input signals must be tied High, the input register must be selected, and the CE and RST input control signals must be tied Low. An example of unused C input recommended settings would be setting C[47:0] = all ones, CREG = 1, CEC = 0, and RSTC = 0.
- These signals are dedicated routing paths internal to the DSP48E2 column. They are not accessible via general routing resources.
- All signals are active High.

Because configuring DSPs is so complicated, Xilinx provides...

```

module DSP48E2.v
//
// Copyright (c) 1995/2017 Xilinx, Inc.
// All Right Reserved.
//
// =====
//
// Vendor      : Xilinx
// Version     : 2017.3
// Description  : Xilinx Unified Simulation
//               48-bit Multi-Functional
// Filename    : DSP48E2.v
//
// =====
timescale 1 ps / 1 ps
celldefine
module DSP48E2 #(
`ifdef XIL_TIMING
parameter LOC = "UNPLACED",
endif
parameter integer ACASCREG = 1,
parameter integer ADREG = 1,
parameter integer ALUMODEREG = 1,
parameter AMULTSEL = "A",
parameter integer AREG = 1,
parameter AUTORESET_PATDET = "NO_RESET",
parameter AUTORESET_PRIORITY = "RESET",
parameter A_INPUT = "DIRECT",
parameter integer BCASCREG = 1,
parameter BMULTSEL = "B",
...
output [29:0] ACOUT,
output [17:0] BCOUT,
output CARRYCASCOUT,
...
input [29:0] A,
input [29:0] ACIN,
input [3:0] ALUNODE,
input [17:0] B,
input [17:0] BCIN,
input [47:0] C,
...
);

// define constants
localparam MODULE_NAME = "DSP48E2";

// Parameter encodings and registers
localparam AMULTSEL_A = 0;
localparam AMULTSEL_AD = 1;
localparam AUTORESET_PATDET_NO_RESET = 0;

`endif

assign ACIN_in = ACIN;
assign ALUNODE_in[0] = (ALUNODE[0] == 1'b1) && (ALUNODE[0] ^
IS_ALUNODE_INVERTED_REG[0]); // rv 0
assign ALUNODE_in[1] = (ALUNODE[1] == 1'b1) && (ALUNODE[1] ^
IS_ALUNODE_INVERTED_REG[1]); // rv 0
assign ALUNODE_in[2] = (ALUNODE[2] == 1'b1) && (ALUNODE[2] ^
IS_ALUNODE_INVERTED_REG[2]); // rv 0
assign ALUNODE_in[3] = (ALUNODE[3] == 1'b1) && (ALUNODE[3] ^
IS_ALUNODE_INVERTED_REG[3]); // rv 0
assign A_in[0] = (A[0] == 1'b1) || A[0]; // rv 1
assign A_in[1] = (A[1] == 1'b1) || A[1]; // rv 1
assign A_in[11] = (A[11] == 1'b1) || A[11]; // rv 1
assign A_in[12] = (A[12] == 1'b1) || A[12]; // rv 1
assign A_in[13] = (A[13] == 1'b1) || A[13]; // rv 1

assign B_in[3] = (B[3] == 1'b1) || B[3]; // rv 1
assign B_in[4] = (B[4] == 1'b1) || B[4]; // rv 1
assign B_in[5] = (B[5] == 1'b1) || B[5]; // rv 1
assign B_in[6] = (B[6] == 1'b1) || B[6]; // rv 1
assign B_in[7] = (B[7] == 1'b1) || B[7]; // rv 1
assign B_in[8] = (B[8] == 1'b1) || B[8]; // rv 1
assign B_in[9] = (B[9] == 1'b1) || B[9]; // rv 1
assign CARRYCASC_in = CARRYCASCIN;
assign CARRYINSEL_in[0] = (CARRYINSEL[0] == 1'b1) &&
CARRYINSEL[0]; // rv 0
assign CARRYINSEL_in[1] = (CARRYINSEL[1] == 1'b1) &&
CARRYINSEL[1]; // rv 0
assign CARRYINSEL_in[2] = (CARRYINSEL[2] == 1'b1) &&
CARRYINSEL[2]; // rv 0
assign CARRYIN_in = (CARRYIN == 1'b1) && (CARRYIN ^
IS_CARRYIN_INVERTED_REG); // rv 0
assign CEAL_in = (CEA1 == 1'b1) && CEA1; // rv 0
assign CEA2_in = (CEA2 == 1'b1) && CEA2; // rv 0
assign CEAD_in = (CEAD == 1'b1) && CEAD; // rv 0
assign CEALUNODE_in = (CEALUNODE == 1'b1) && CEALUNODE; // rv 0
assign CEB1_in = (CEB1 == 1'b1) && CEB1; // rv 0
assign CEB2_in = (CEB2 == 1'b1) && CEB2; // rv 0

assign CECARRYIN_in = (CECARRYIN == 1'b1) && CECARRYIN; // rv 0
assign CECTRL_in = (CECTRL == 1'b1) && CECTRL; // rv 0
assign CEC_in = (CEC == 1'b1) && CEC; // rv 0
assign CED_in = (CED == 1'b1) && CED; // rv 0
assign CEINMODE_in = CEINMODE;
assign CER_in = (CER == 1'b1) && CER; // rv 0
assign CEP_in = (CEP == 1'b1) && CEP; // rv 0
assign CLK_in = (CLK == 1'b1) && (CLK ^ IS_CLK_INVERTED_REG);
// rv 0
assign C_in[0] = (C[0] == 1'b1) || C[0]; // rv 1
assign C_in[10] = (C[10] == 1'b1) || C[10]; // rv 1
a
assign D_in[1] = (D[1] == 1'b1) && D[1]; // rv 0
assign D_in[20] = (D[20] == 1'b1) && D[20]; // rv 0
assign D_in[21] = (D[21] == 1'b1) && D[21]; // rv 0
assign D_in[22] = (D[22] == 1'b1) && D[22]; // rv 0
assign D_in[23] = (D[23] == 1'b1) && D[23]; // rv 0
assign D_in[24] = (D[24] == 1'b1) && D[24]; // rv 0
assign D_in[25] = (D[25] == 1'b1) && D[25]; // rv 0
assign D_in[26] = (D[26] == 1'b1) && D[26]; // rv 0
assign D_in[2] = (D[2] == 1'b1) && D[2]; // rv 0
assign D_in[3] = (D[3] == 1'b1) && D[3]; // rv 0
assign D_in[4] = (D[4] == 1'b1) && D[4]; // rv 0
assign D_in[5] = (D[5] == 1'b1) && D[5]; // rv 0
assign D_in[6] = (D[6] == 1'b1) && D[6]; // rv 0
assign D_in[7] = (D[7] == 1'b1) && D[7]; // rv 0
assign D_in[8] = (D[8] == 1'b1) && D[8]; // rv 0
assign D_in[9] = (D[9] == 1'b1) && D[9]; // rv 0
assign INMODE_in[0] = (INMODE[0] == 1'b1) && (INMODE[0] ^
IS_INMODE_INVERTED_REG[0]); // rv 0
assign INMODE_in[1] = (INMODE[1] == 1'b1) && (INMODE[1] ^
IS_INMODE_INVERTED_REG[1]); // rv 0
assign INMODE_in[2] = (INMODE[2] == 1'b1) && (INMODE[2] ^
IS_INMODE_INVERTED_REG[2]); // rv 0
assign INMODE_in[3] = (INMODE[3] == 1'b1) && (INMODE[3] ^
IS_INMODE_INVERTED_REG[3]); // rv 0
assign INMODE_in[4] = (INMODE[4] == 1'b1) && (INMODE[4] ^
IS_INMODE_INVERTED_REG[4]); // rv 0
assign MULTSIGNIN_in = MULTSIGNIN;
assign OPHODE_in[0] = (OPHODE[0] == 1'b1) && (OPHODE[0] ^
IS_OPHODE_INVERTED_REG[0]); // rv 0
assign OPHODE_in[1] = (OPHODE[1] == 1'b1) && (OPHODE[1] ^
IS_OPHODE_INVERTED_REG[1]); // rv 0
assign OPHODE_in[2] = (OPHODE[2] == 1'b1) && (OPHODE[2] ^
IS_OPHODE_INVERTED_REG[2]); // rv 0
assign OPHODE_in[3] = (OPHODE[3] == 1'b1) && (OPHODE[3] ^
IS_OPHODE_INVERTED_REG[3]); // rv 0
...

```

a 1500 line simulation model of the DSP, which hardware designers use to validate their designs.

But this is useful to us as well, because

[build] ...

module DSP48E2.v

```

// Copyright (c) 1995/2017 Xilinx, Inc.
// All Right Reserved.
//
// Vendor      : Xilinx
// Version     : 2017.3
// Description  : Xilinx Unified Simulation
//              : 48-bit Multi-Functional
// Filename    : DSP48E2.v
//
timescale 1 ps / 1 ps

module DSP48E2
    parameter integer ALUMODESEL = 1,
    parameter integer AMULTSEL = "A",
    parameter integer AREG = 1,
    parameter AUTORESET_PATDET = "NO_RESET",
    parameter AUTORESET_PRIORITY = "RESET",
    parameter A_INPUT = "DIRECT",
    parameter integer BCASCSEL = 1,
    parameter BMULTSEL = "B",
    output [29:0] ACOUT,
    output [17:0] BCOUT,
    output CARRYCASCOUT,
    ...
    input [29:0] A,
    input [29:0] ACIN,
    input [3:0] ALUMODE,
    input [17:0] B,
    input [17:0] BCIN,
    input [47:0] C,
    ...
endmodule

// define constants
localparam MODULE_NAME = "DSP48E2";

// Parameter encodings and registers
localparam AMULTSEL_A = 0;
localparam AMULTSEL_AD = 1;
localparam AUTORESET_PATDET_NO_RESET = 0;

`endif

assign ACIN_in = ACIN;
assign ALUMODE_in[0] = (ALUMODE[0] == 1'b1) && (ALUMODE[0] ^
IS_ALUMODE_INVERTED_REG[0]); // rv 0
assign ALUMODE_in[1] = (ALUMODE[1] == 1'b1) && (ALUMODE[1] ^
IS_ALUMODE_INVERTED_REG[1]); // rv 0
assign ALUMODE_in[2] = (ALUMODE[2] == 1'b1) && (ALUMODE[2] ^
IS_ALUMODE_INVERTED_REG[2]); // rv 0
assign ALUMODE_in[3] = (ALUMODE[3] == 1'b1) && (ALUMODE[3] ^
IS_ALUMODE_INVERTED_REG[3]); // rv 0

assign CECARRYIN_in = (CECARRYIN == 1'b1) && CECARRYIN; // rv 0
assign CECTRL_in = (CECTRL == 1'b1) && CECTRL; // rv 0
assign CEC_in = (CEC == 1'b1) && CEC; // rv 0
assign CED_in = (CED == 1'b1) && CED; // rv 0
assign CEINMODE_in = CEINMODE;
assign CEN_in = (CEN == 1'b1) && CEN; // rv 0
assign CEP_in = (CEP == 1'b1) && CEP; // rv 0
assign CLK_in = (CLK == 1'b1) && (CLK ^ IS_CLK_INVERTED_REG);
// rv 0
assign C_in[0] = (C[0] == 1'b1) || C[0]; // rv 1
assign C_in[10] = (C[10] == 1'b1) || C[10]; // rv 1
a
assign D_in[1] = (D[1] == 1'b1) && D[1]; // rv 0
assign D_in[20] = (D[20] == 1'b1) && D[20]; // rv 0
assign D_in[21] = (D[21] == 1'b1) && D[21]; // rv 0
assign D_in[22] = (D[22] == 1'b1) && D[22]; // rv 0
assign D_in[23] = (D[23] == 1'b1) && D[23]; // rv 0
assign D_in[24] = (D[24] == 1'b1) && D[24]; // rv 0
assign D_in[25] = (D[25] == 1'b1) && D[25]; // rv 0
assign D_in[26] = (D[26] == 1'b1) && D[26]; // rv 0
assign D_in[27] = (D[27] == 1'b1) && D[27]; // rv 0

assign B_in[5] = (B[5] == 1'b1) || B[5]; // rv 1
assign B_in[6] = (B[6] == 1'b1) || B[6]; // rv 1
assign B_in[7] = (B[7] == 1'b1) || B[7]; // rv 1
assign B_in[8] = (B[8] == 1'b1) || B[8]; // rv 1
assign B_in[9] = (B[9] == 1'b1) || B[9]; // rv 1
assign CARRYCASCIN_in = CARRYCASCIN;
assign CARRYINSEL_in[0] = (CARRYINSEL[0] == 1'b1) &&
CARRYINSEL[0]; // rv 0
assign CARRYINSEL_in[1] = (CARRYINSEL[1] == 1'b1) &&
CARRYINSEL[1]; // rv 0
assign CARRYINSEL_in[2] = (CARRYINSEL[2] == 1'b1) &&
CARRYINSEL[2]; // rv 0
assign CARRYIN_in = (CARRYIN == 1'b1) && (CARRYIN ^
IS_CARRYIN_INVERTED_REG); // rv 0
assign CEAL_in = (CEAL == 1'b1) && CEAL; // rv 0
assign CEAL_in = (CEAL == 1'b1) && CEAL; // rv 0
assign CEAD_in = (CEAD == 1'b1) && CEAD; // rv 0
assign CEALUMODE_in = (CEALUMODE == 1'b1) && CEALUMODE; // rv 0
assign CEB1_in = (CEB1 == 1'b1) && CEB1; // rv 0
assign CEB2_in = (CEB2 == 1'b1) && CEB2; // rv 0

IS_INMODE_INVERTED_REG[0]; // rv 0
assign INMODE_in[1] = (INMODE[1] == 1'b1) && (INMODE[1] ^
IS_INMODE_INVERTED_REG[1]); // rv 0
assign INMODE_in[2] = (INMODE[2] == 1'b1) && (INMODE[2] ^
IS_INMODE_INVERTED_REG[2]); // rv 0
assign INMODE_in[3] = (INMODE[3] == 1'b1) && (INMODE[3] ^
IS_INMODE_INVERTED_REG[3]); // rv 0
assign INMODE_in[4] = (INMODE[4] == 1'b1) && (INMODE[4] ^
IS_INMODE_INVERTED_REG[4]); // rv 0
assign MULTSIGNIN_in = MULTSIGNIN;
assign OPNODE_in[0] = (OPNODE[0] == 1'b1) && (OPNODE[0] ^
IS_OPNODE_INVERTED_REG[0]); // rv 0
assign OPNODE_in[1] = (OPNODE[1] == 1'b1) && (OPNODE[1] ^
IS_OPNODE_INVERTED_REG[1]); // rv 0
assign OPNODE_in[2] = (OPNODE[2] == 1'b1) && (OPNODE[2] ^
IS_OPNODE_INVERTED_REG[2]); // rv 0
assign OPNODE_in[3] = (OPNODE[3] == 1'b1) && (OPNODE[3] ^
IS_OPNODE_INVERTED_REG[3]); // rv 0

...

```

Simulation models provide the formal semantics of behaviors and constraints necessary for automated reasoning!

Generating Compilers → Why Now? → Case Study: Lakeroad → Call to Action

```
module DSP48E2 #(
  par
  par
  par
  par
  par
  par
  par
  par
  ...
)(
  out
  out
  out
  ...
  inp
  inp
  inp
  inp
  inp
  ...
);
```

Insight #1: configuring DSPs and other complex primitives is similar to writing a program, so use *program synthesis*.

Insight #2: we can extract the semantics necessary for automated reasoning directly from simulation models.

This leads us to our second insight, which is that we can ...

Lakeroad: a hardware synthesis tool utilizing program synthesis and semantics extracted from simulation models to target complex, programmable FPGA primitives.

Using these two insights, we build Lakeroad, which is...

| Workload | Signed? | # Stages | Yosys | SOTA | Lakeroad |
|--------------------------|---------|----------|---------------|---------------|----------|
| $((d + a) * b) \mid c$ | X | 1 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) \mid c$ | ✓ | 2 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) \wedge c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) \& c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) \wedge c$ | X | 2 | 1 DSP, 18 LUT | 1 DSP, 9 LUT | 1 DSP |

Initial results with Lakeroad suggest that Lakeroad is able to find mappings that other tools do not, successfully mapping designs onto [build] a single DSP when other tools [build] fail to do so.

| Workload | Signed? | # Stages | Yosys | SOTA | Lakeroad |
|------------------------|---------|----------|---------------|---------------|----------|
| $((d + a) * b) \mid c$ | X | 1 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) \mid c$ | ✓ | 2 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) ^ c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) \& c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) ^ c$ | X | 2 | 1 DSP, 18 LUT | 1 DSP, 9 LUT | 1 DSP |

| Workload | Signed? | # Stages | Yosys | SOTA | Lakeroad |
|------------------------|---------|----------|---------------|---------------|----------|
| $((d + a) * b) \mid c$ | X | 1 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) \mid c$ | ✓ | 2 | 1 DSP, 20 LUT | 1 DSP, 10 LUT | 1 DSP |
| $((d - a) * b) ^ c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) \& c$ | ✓ | 3 | 1 DSP, 22 LUT | 2 DSP, 11 LUT | 1 DSP |
| $((d + a) * b) ^ c$ | X | 2 | 1 DSP, 18 LUT | 1 DSP, 9 LUT | 1 DSP |

How does Lakeroad support the thesis of this talk?

We claim that

[build] ...

Lakeroad exemplifies this directly, by

[build] automatically enabling compilation...

Furthermore, we claim that

[build] ...

Which Lakeroad demonstrates by showing that

[build] automated methods are...

**Compilers should be generated from formal
models of hardware.**

Compilers should be generated from formal models of hardware.

Lakeroad automatically enables compilation to FPGA primitives, given the simulation models of those primitives.

Compilers should be generated from formal models of hardware.

Lakeroad automatically enables compilation to FPGA primitives, given the simulation models of those primitives.

With the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make this a reality.

Compilers should be generated from formal models of hardware.

Lakeroad automatically enables compilation to FPGA primitives, given the simulation models of those primitives.

With the growing diversity of hardware and the rapid improvement of automated reasoning, now is the time to make this a reality.

Lakeroad demonstrates that automated methods are now powerful enough address gaps in existing state-of-the-art tools.

Generating Compilers → Why Now? → **Case Study: Lakeroad** → Call to Action

Now finally,

Generating Compilers → Why Now? → Case Study: Lakeroad → **Call to Action**

Let's conclude with a call to action.

The Hardware Lottery

Sara Hooker

Google Research, Brain Team
shooker@google.com

2020

In conclusion, I want to end where we began.

I believe the hardware lottery is a direct challenge to our community.

If there's one message I want to leave you with today, it's this:

[build] It's on on all of us in this room to fight against the hardware lottery, by making sure that practitioners in all fields have the hardware and compilers they need to advance their research.

What I've proposed today—automatically generating compilers from formal models of hardware—is just one possible solution. Whether or not you agree with the solution, I hope you'll agree with the larger goal of ending the hardware lottery once and for all.

The Hardware Lottery

Sara Hooker

Google Research, Brain Team
shooker@google.com

2020

It's on on all of us to fight against the hardware lottery, by making sure that practitioners in all fields have the hardware and compilers they need to advance their research.

Thank you!



Thank you, everybody!