


HOW TO GIVE TALKS THAT PEOPLE CAN FOLLOW



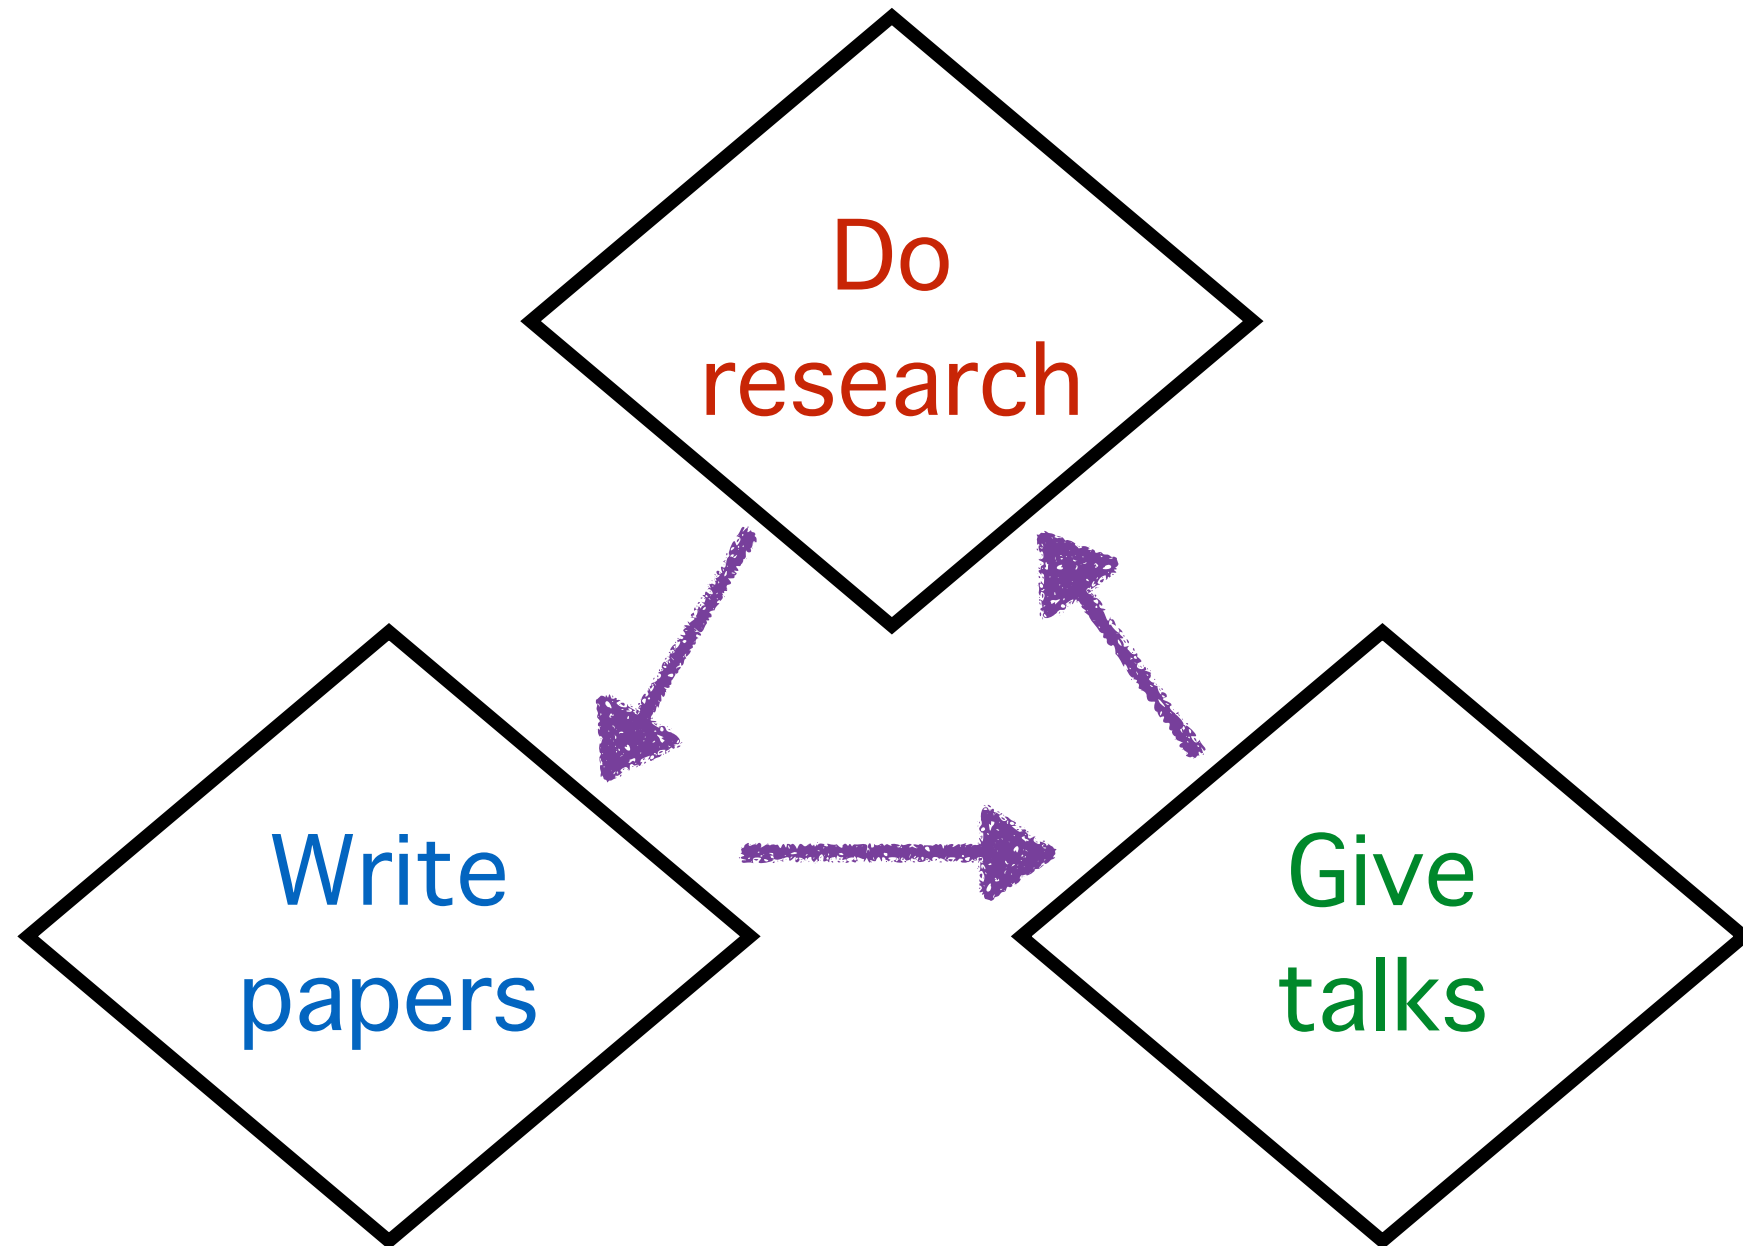
Derek Dreyer
Max Planck Institute for Software Systems

PLMW@PLDI 2021

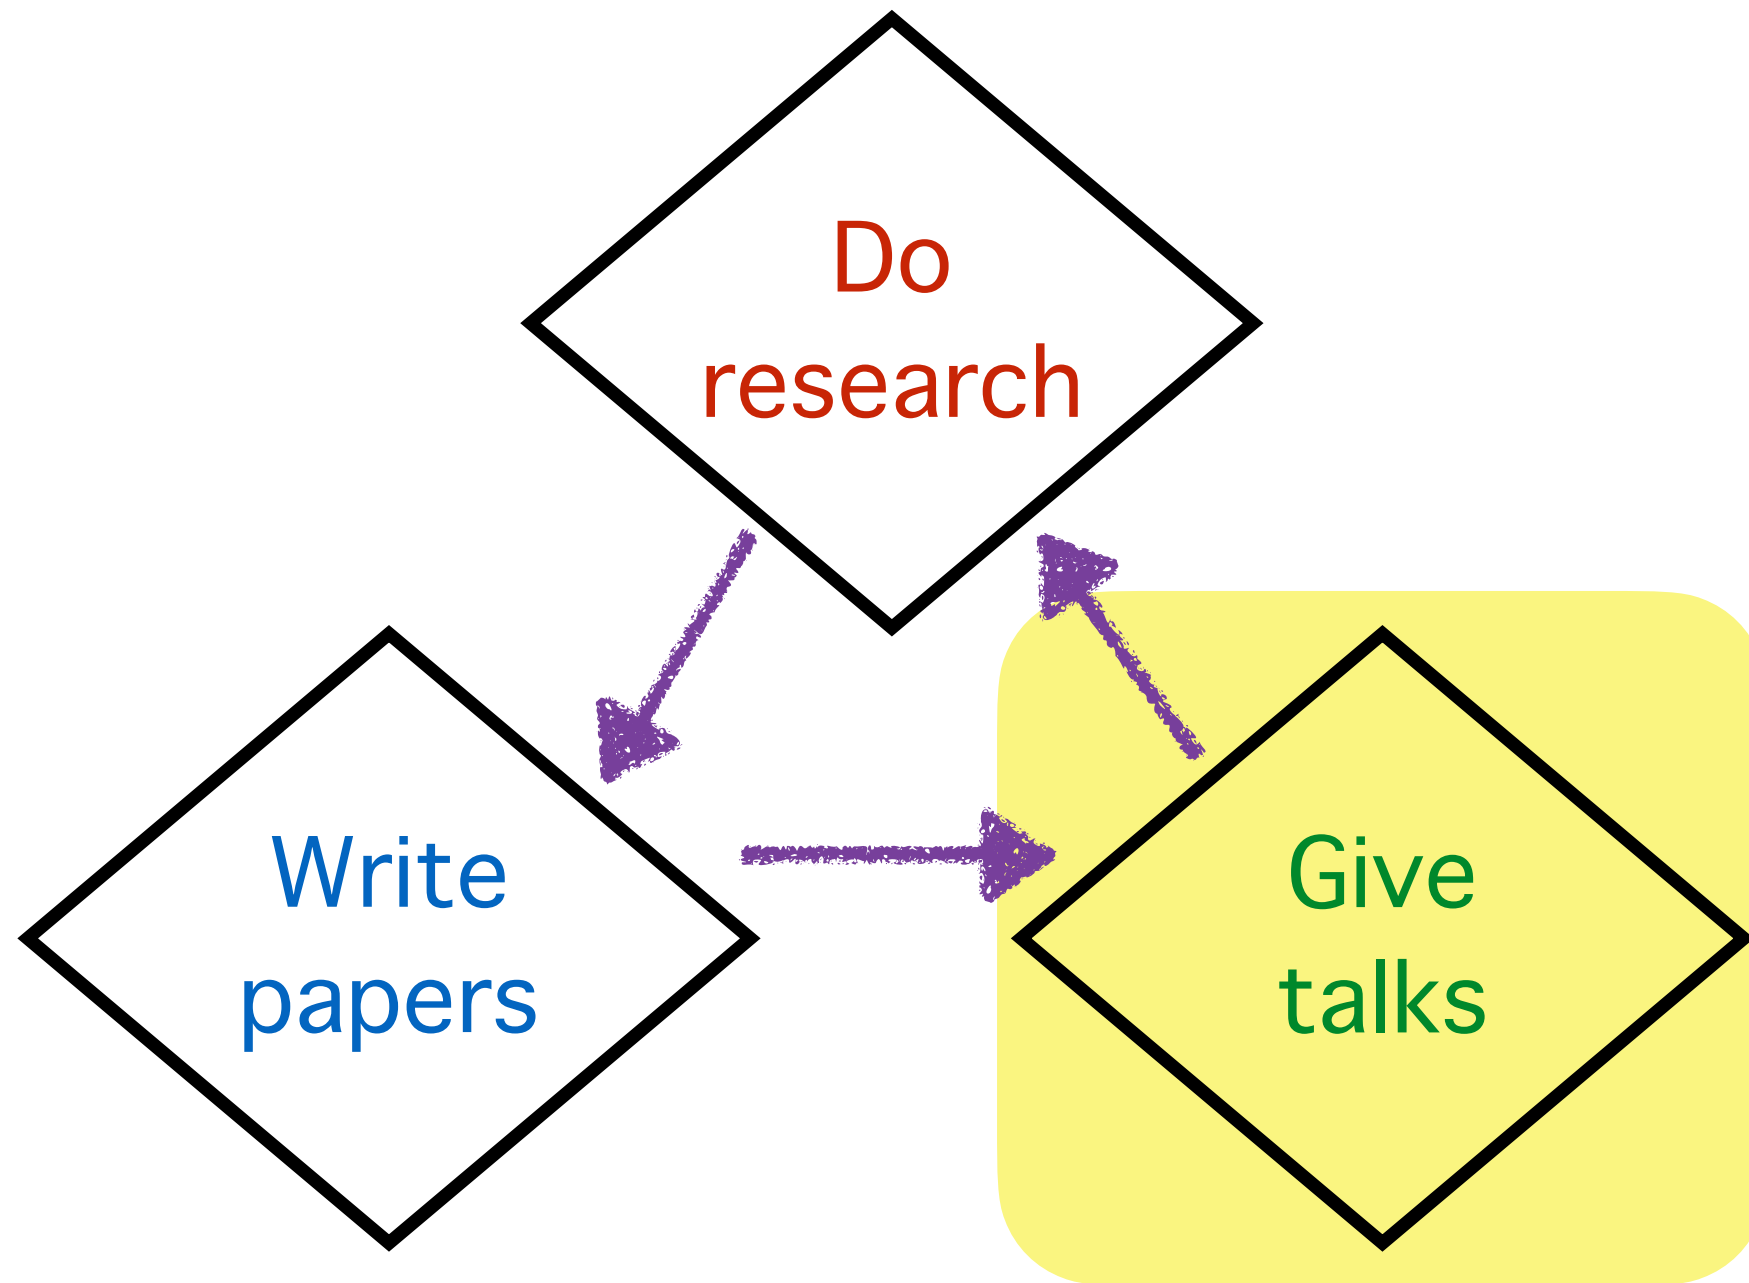
My job as a researcher



My job as a researcher



My job as a researcher



My job as a researcher

Talk developed jointly with
Rose Hoberman
@ MPI-SWS



Write
papers

Give
talks








Entertain your audience!

- **Simon Peyton Jones.** How to give a great research talk. (MSR Summer School, 2016)
 - “Your mission is to **wake them up!**”
 - “Your most potent weapon, by far, is **your enthusiasm!**”
- **John Hughes.** Unaccustomed as I am to public speaking. (PLMW, 2016)
 - “**Put on a show!**”



Entertain your audience!

- **Simon Peyton Jones.** How to give a great research talk. (MSR Summer School, 2016)



Good advice, but I don't know how to teach people to be entertaining...

- **John Hughes.** Unaccustomed as I am to public speaking. (PLMW, 2016)



- “**Put on a show!**”

What is your main goal in
giving a conference talk?

What is your main goal in
giving a conference talk?

Get people to read your paper?

What is your main goal in
giving a conference talk?

Get people to read your paper?

No! Talk \neq Paper

What is your main goal in
giving a conference talk?

**Give people positive feelings
about you and your work!**

Extreme example:
My ICFP'15 talk

Pilsner:

*A Compositionally Verified Compiler for a
Higher-Order Imperative Language*



Georg Neis, Chung-Kil Hur,
Jan-Oliver Kaiser, Craig McLaughlin,
Derek Dreyer, Viktor Vafeiadis

**MPI-SWS (Germany),
Seoul National University,
University of Glasgow**

**ICFP 2015
Vancouver**

Our Contributions

Parametric Inter-Language Simulations (PILS):

- New way to define semantics preservation
- Modular, flexible, *and* transitive

Pilsner:

- The *first* compositionally verified multi-pass compiler for an ML-like language
- Verified using PILS in Coq!

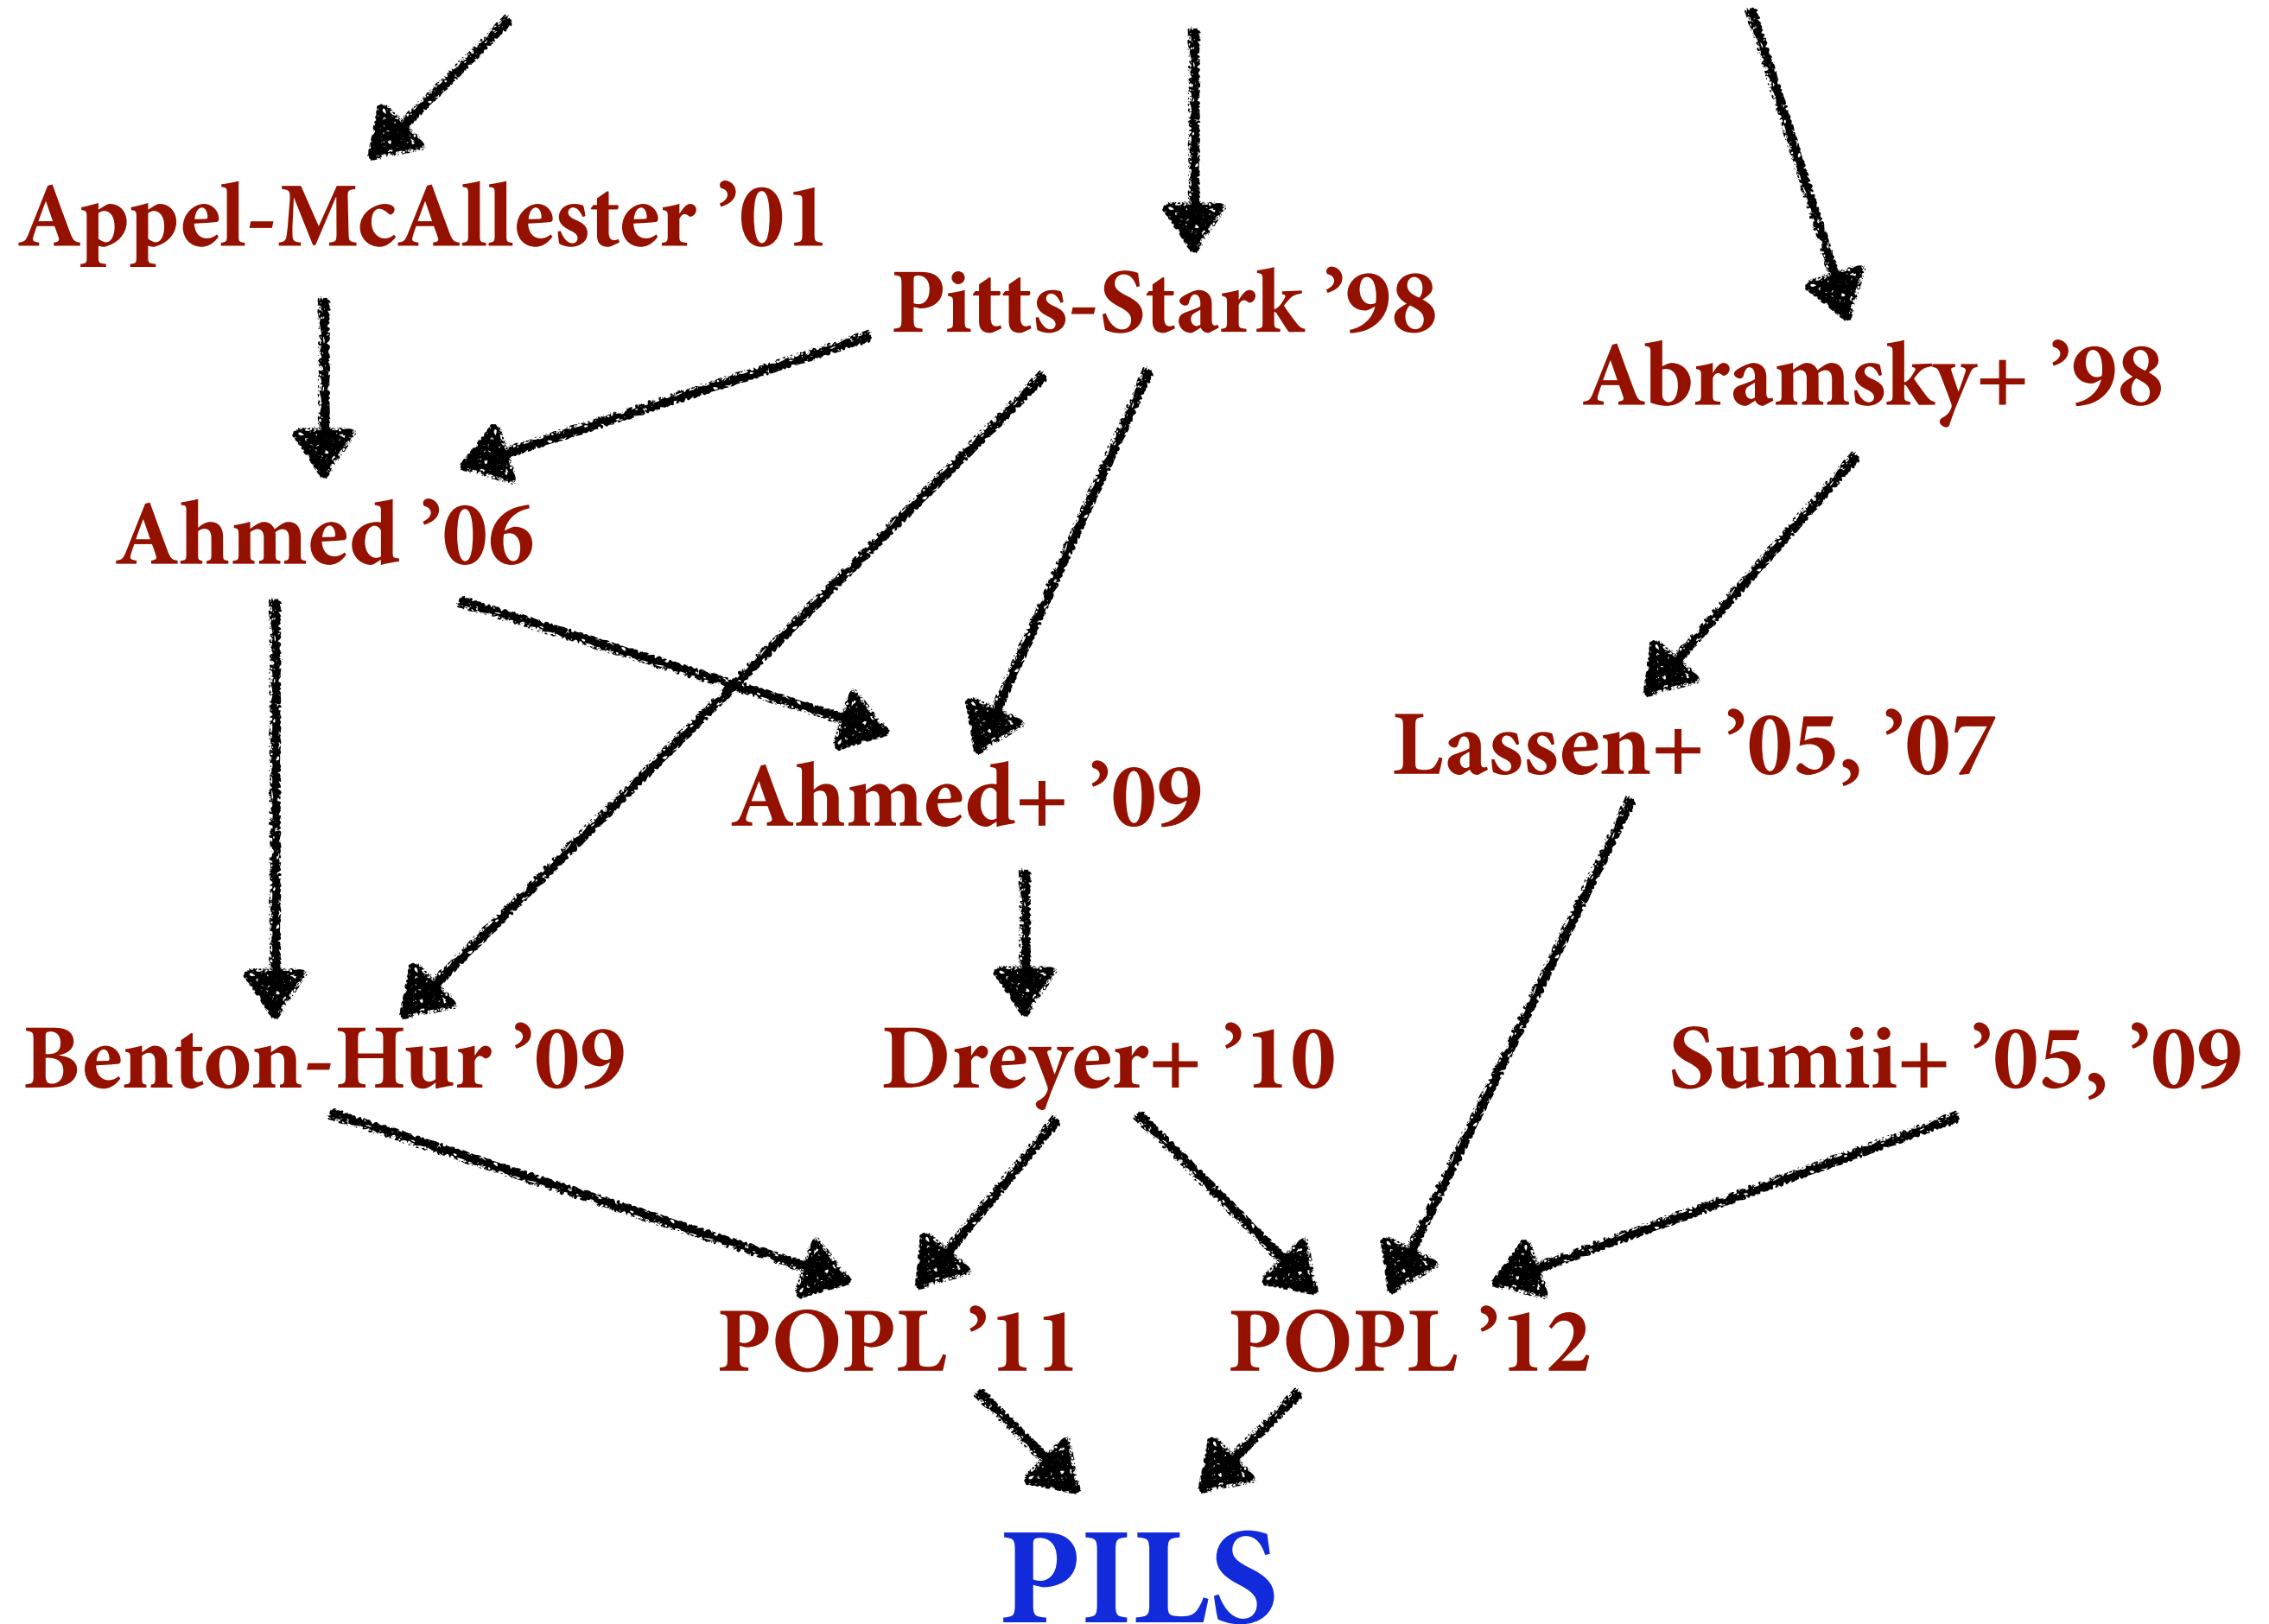
Our Contributions

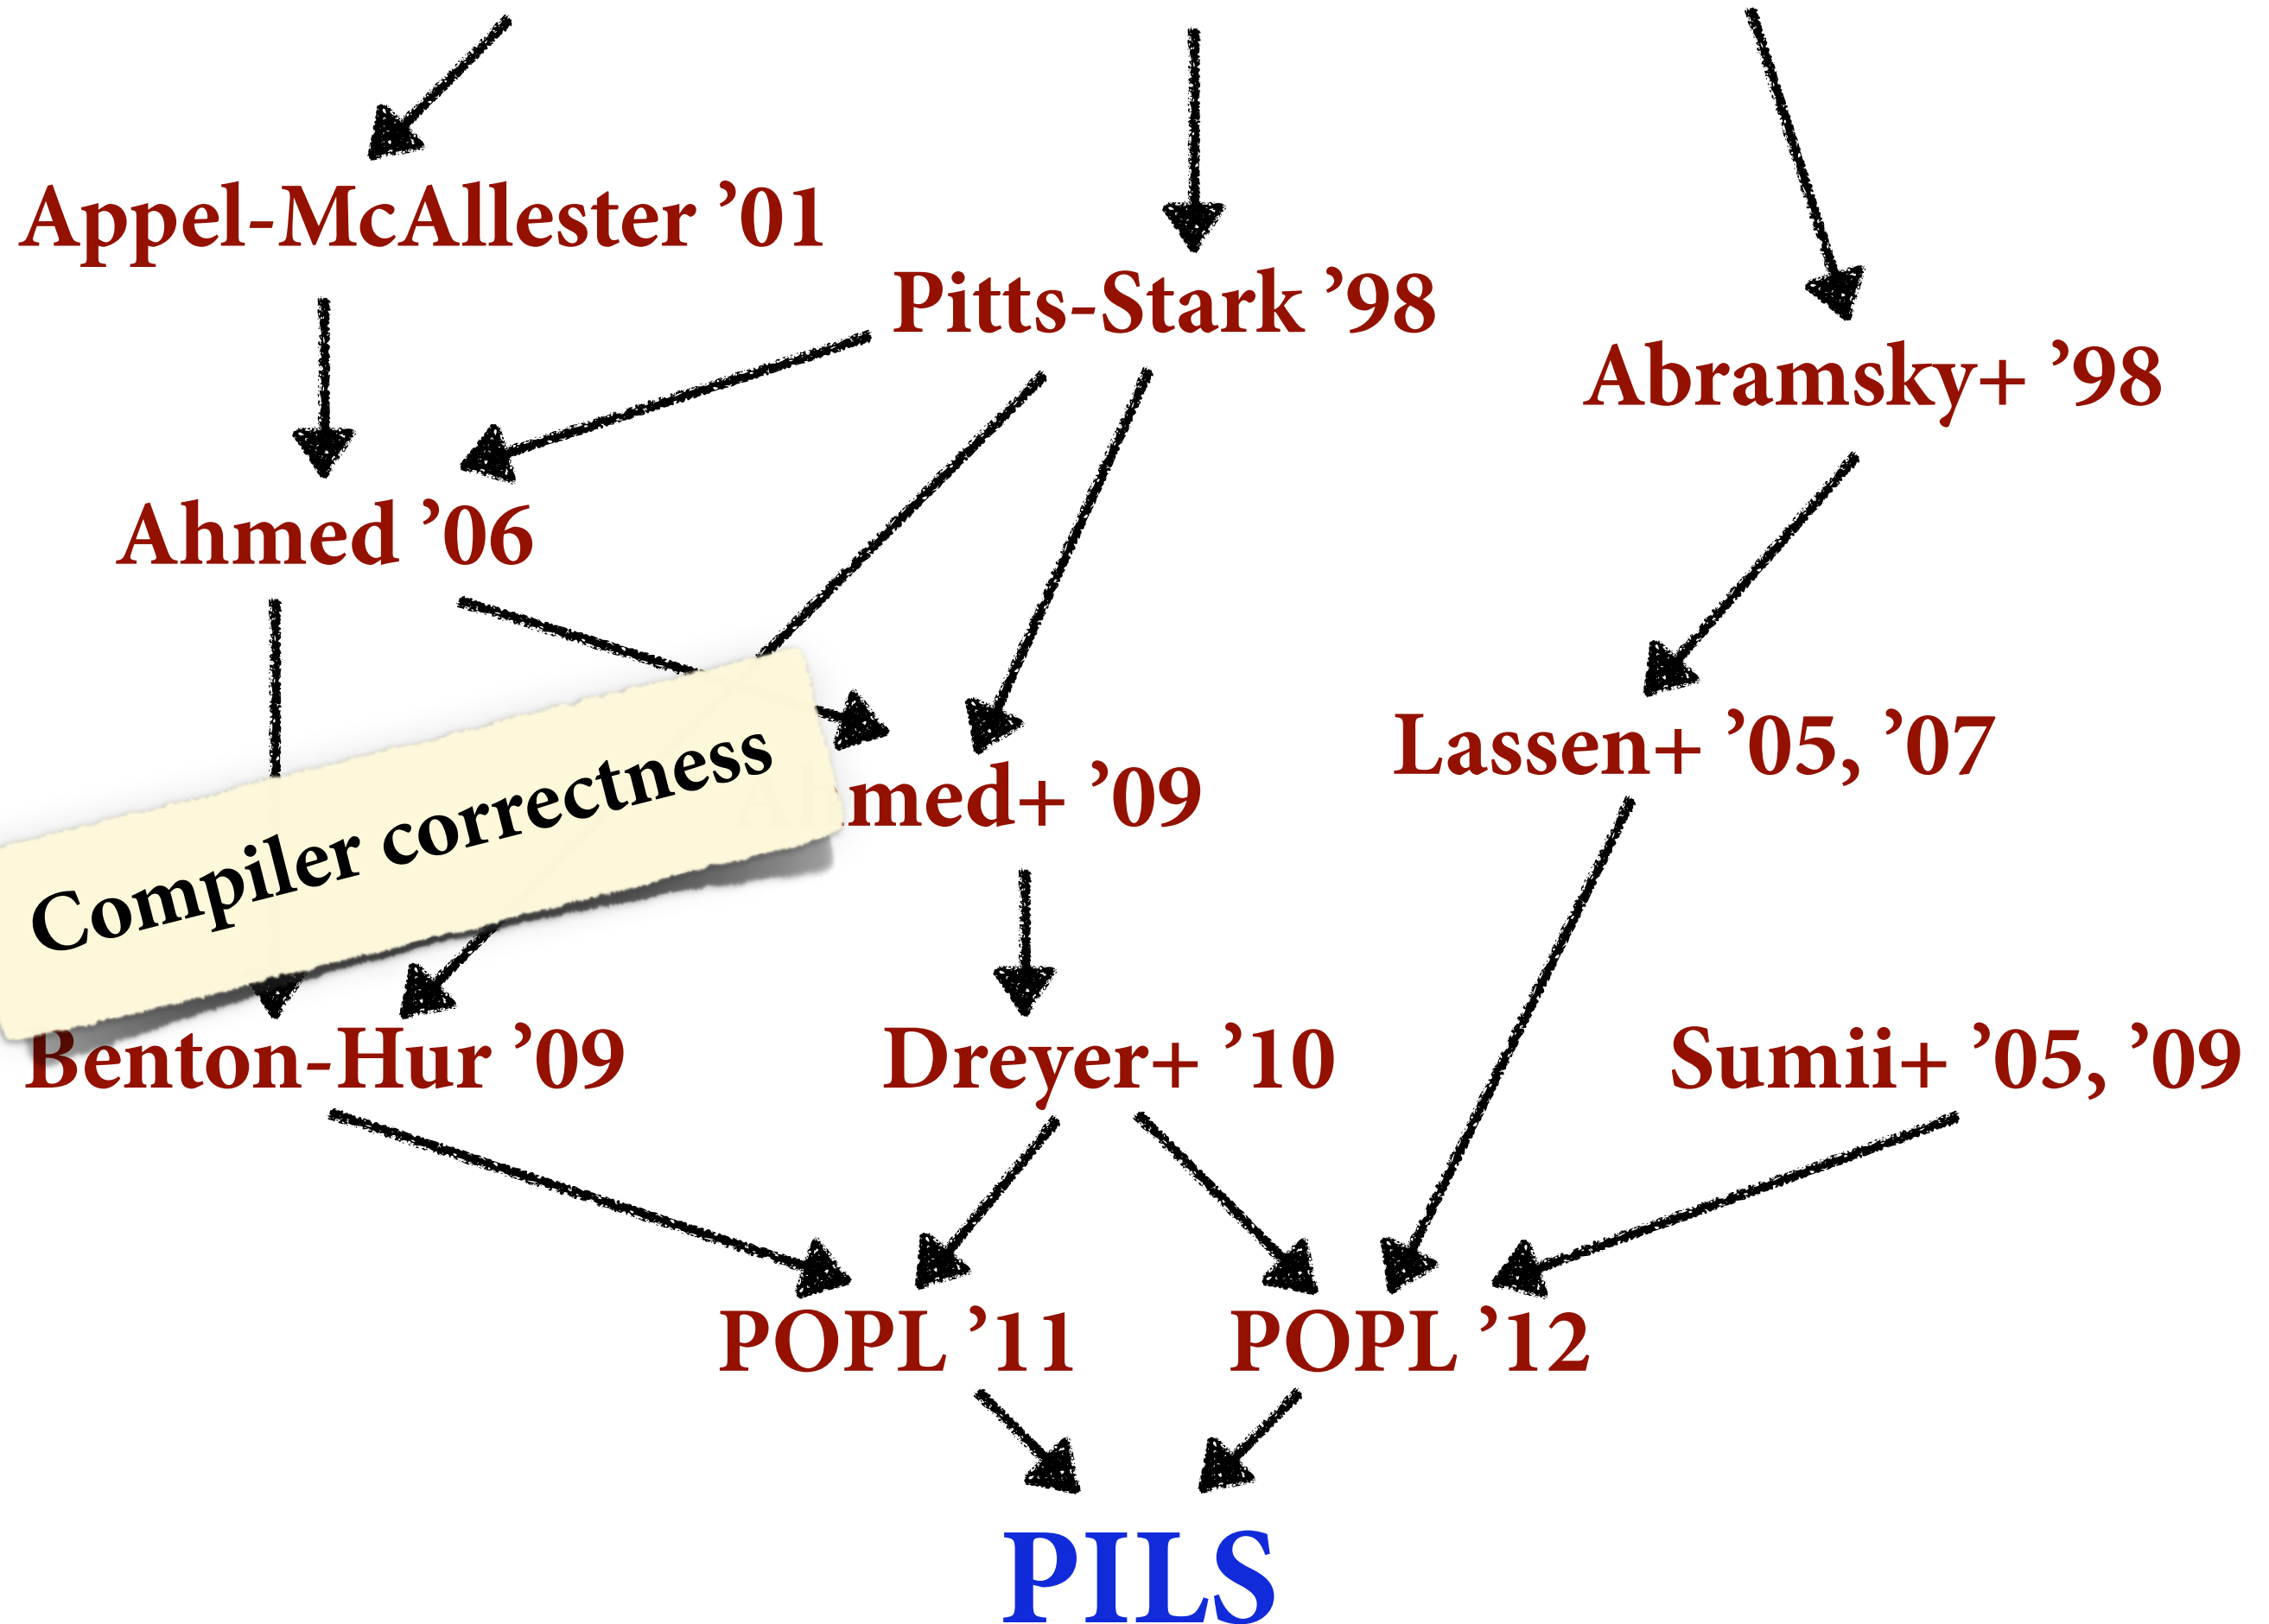
Parametric Inter-Language Simulations (PILS):

- New way to define semantics preservation
- Modular, flexible, *and* transitive

Pilsner:

- The *first* compositionally verified multi-pass compiler for an ML-like language
- Verified using PILS in Coq!





Appel-McAllester '01

Step-indexing

Pitts-Stark '98

Abramsky+ '98

Ahmed '06

Compiler correctness

med+ '09

Lassen+ '05, '07

Benton-Hur '09

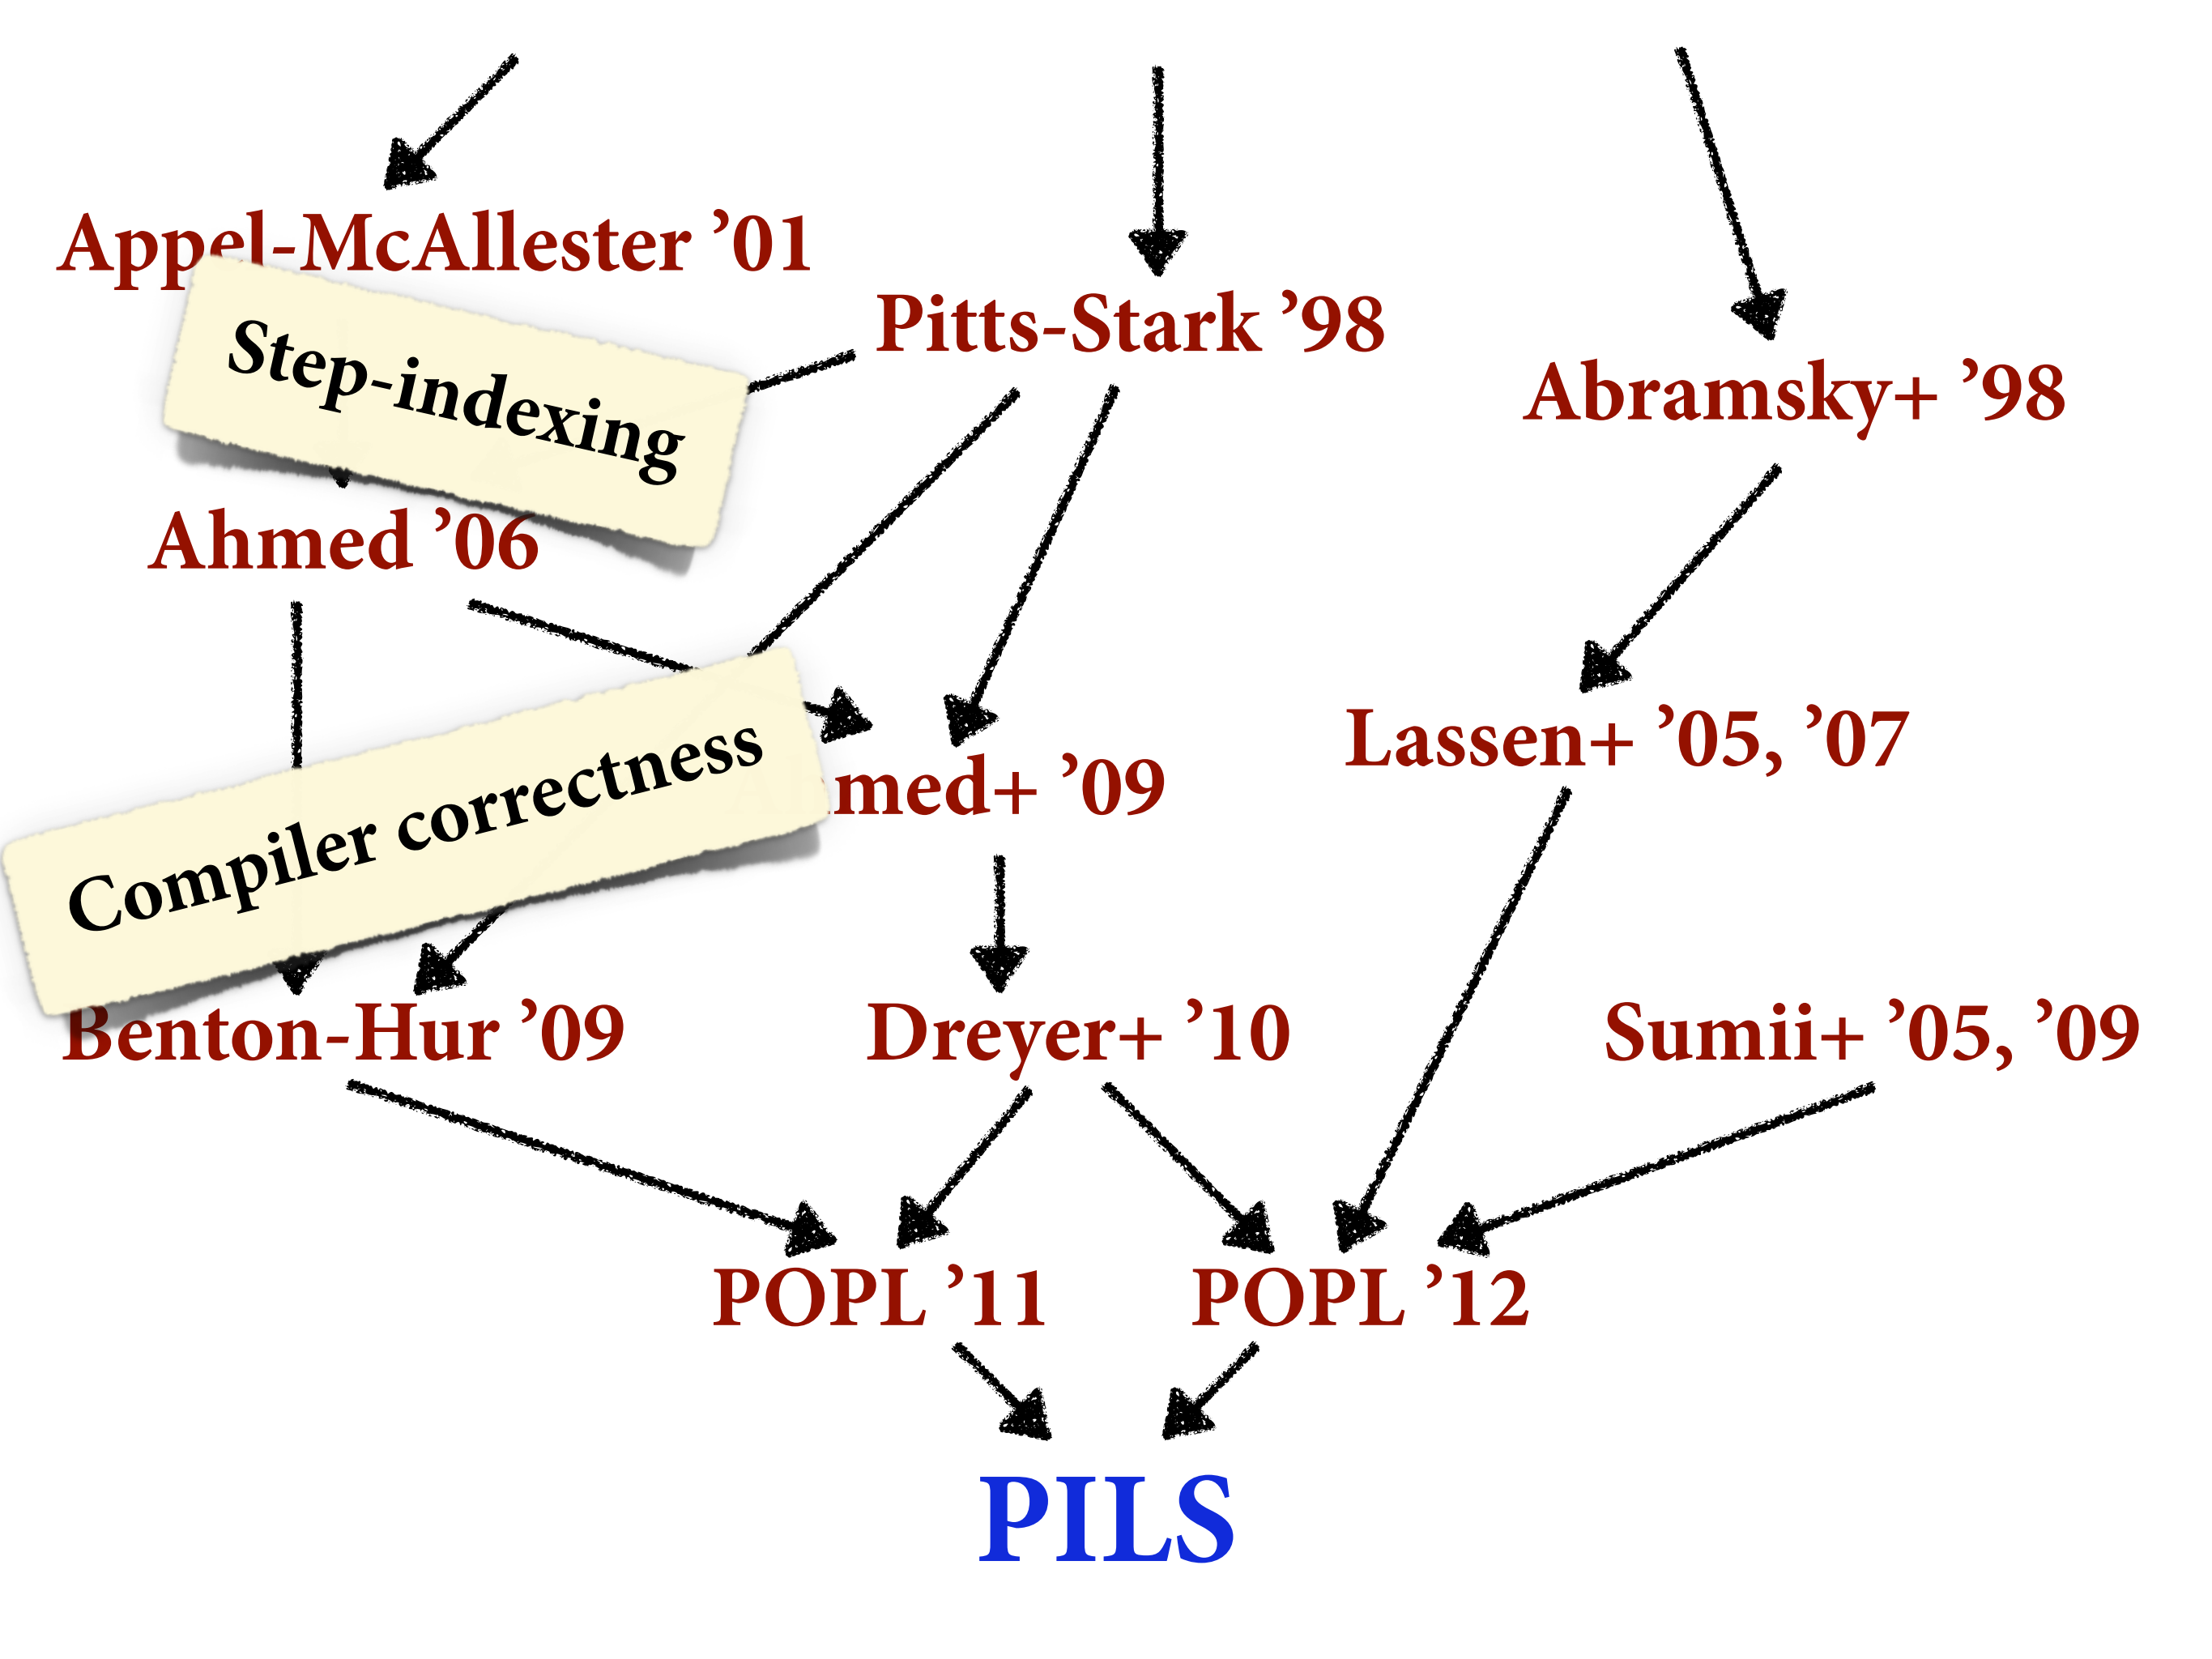
Dreyer+ '10

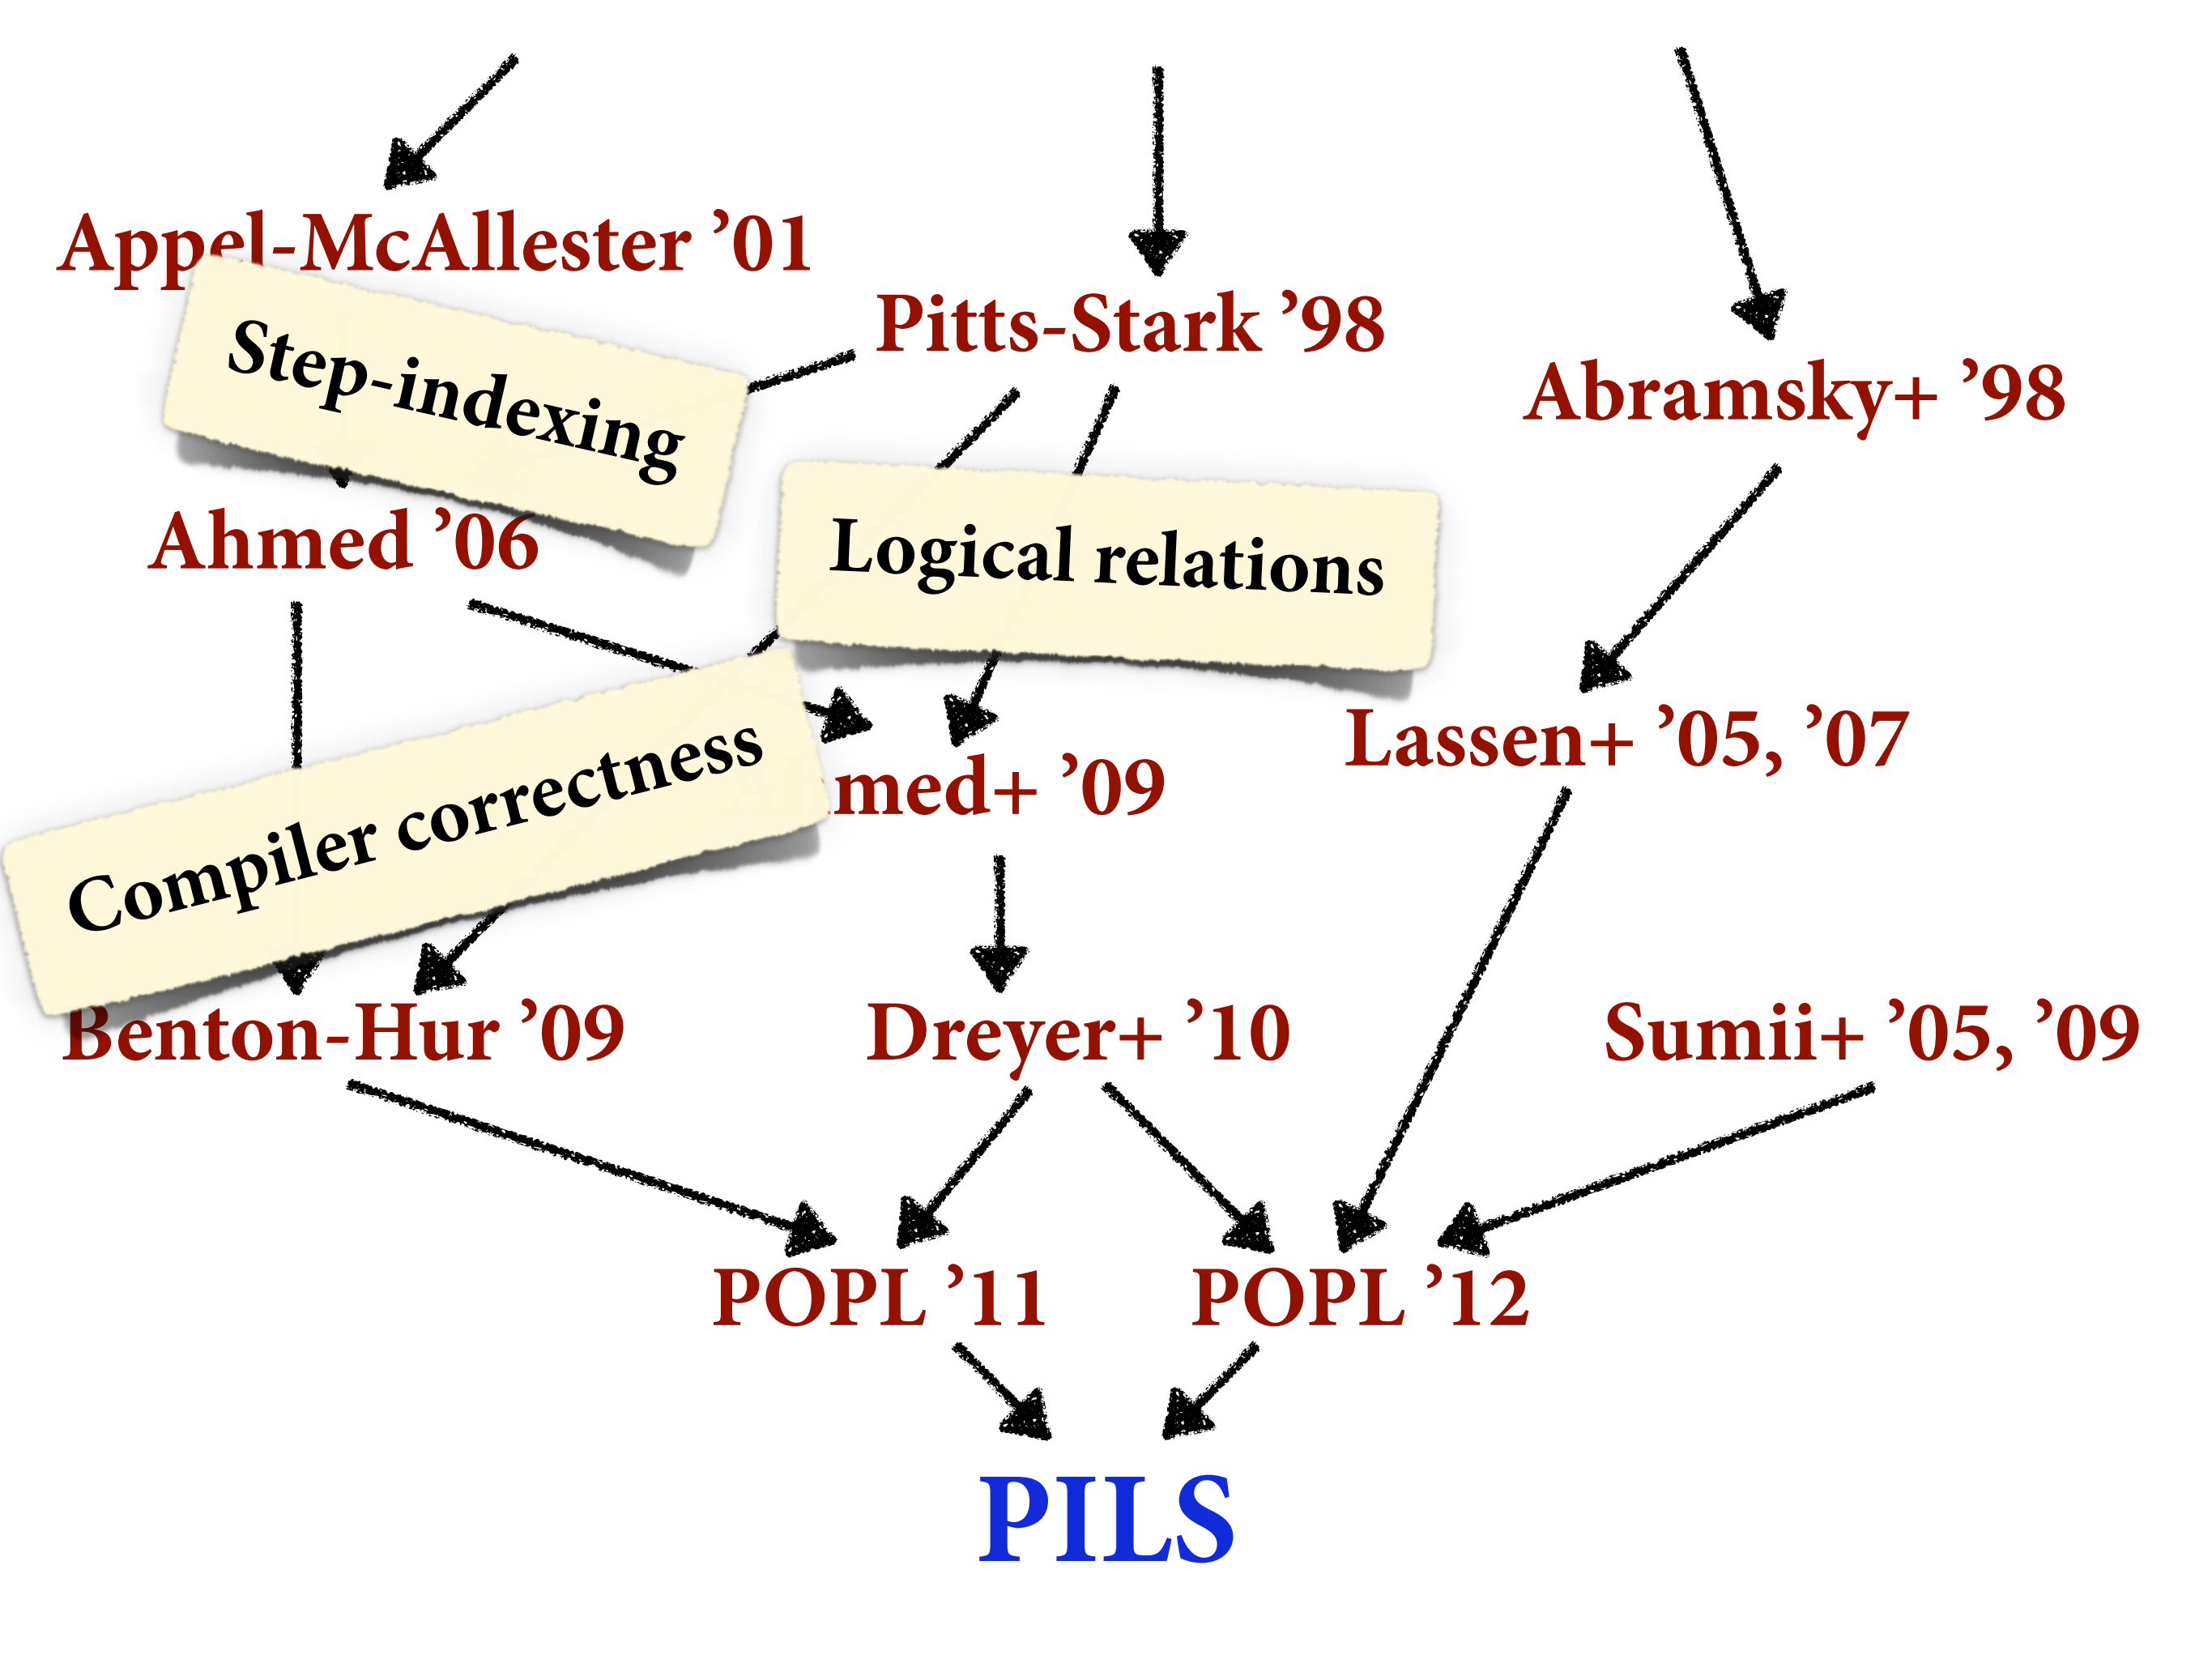
Sumii+ '05, '09

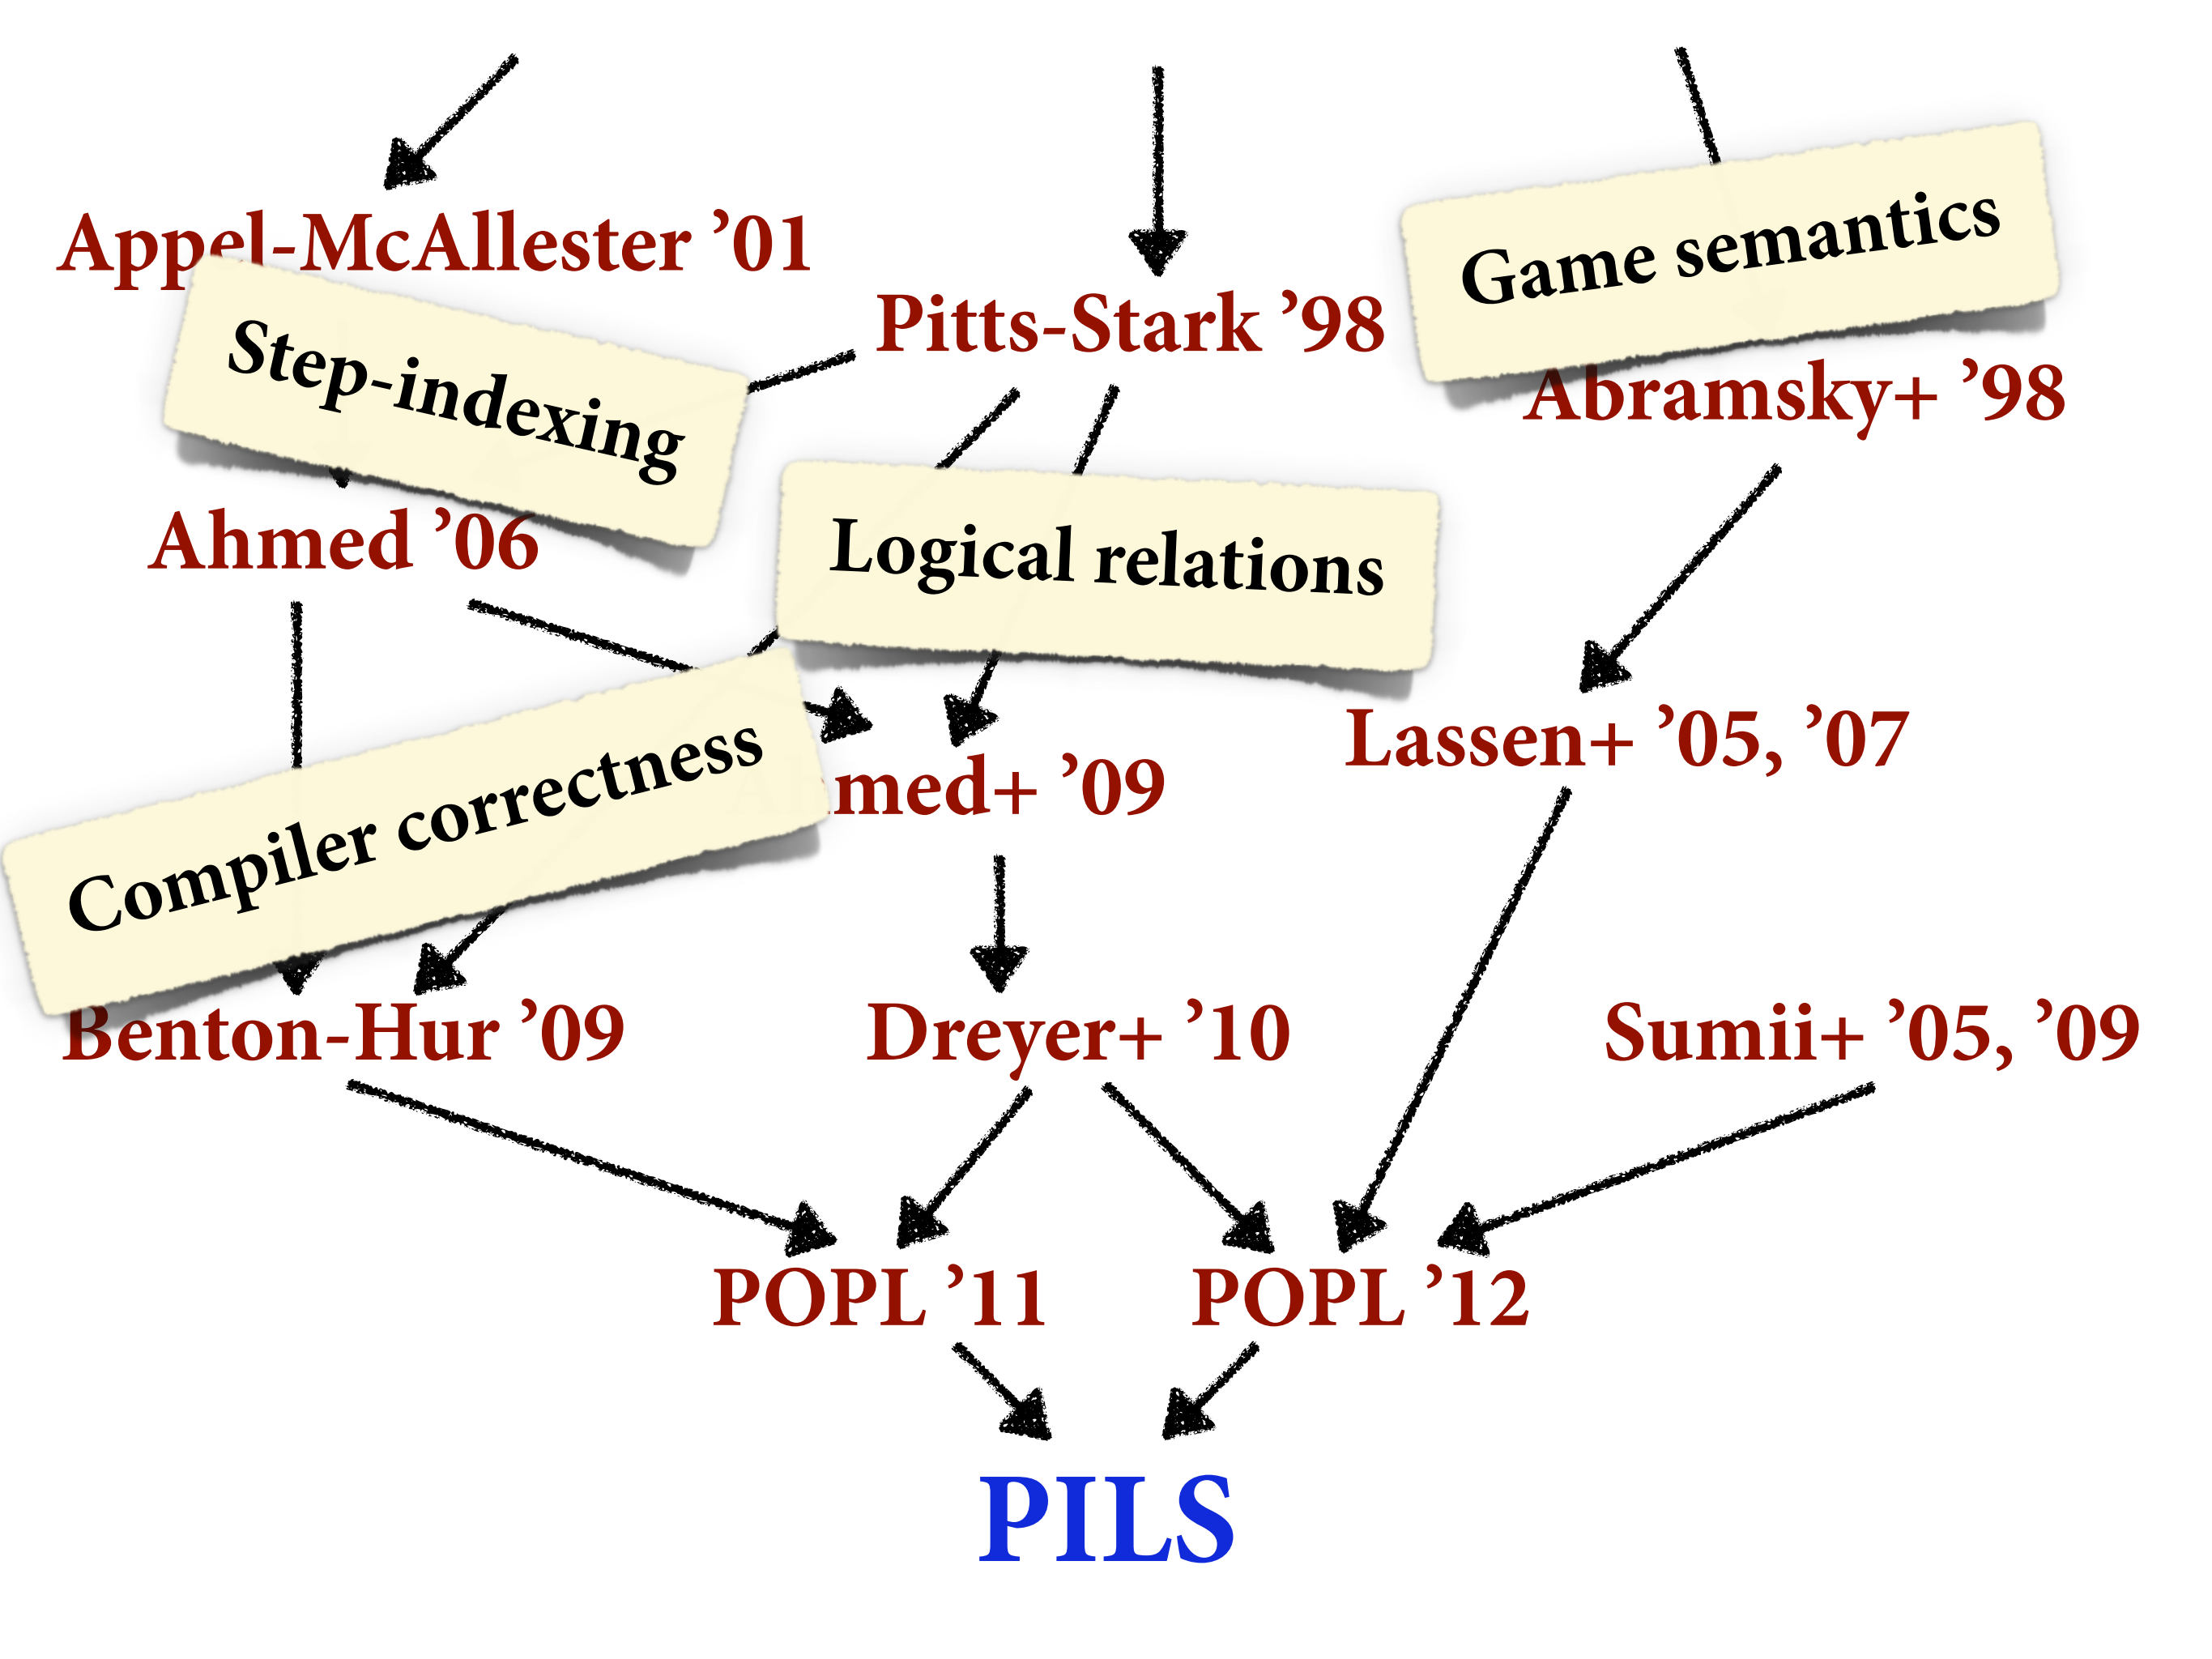
POPL '11

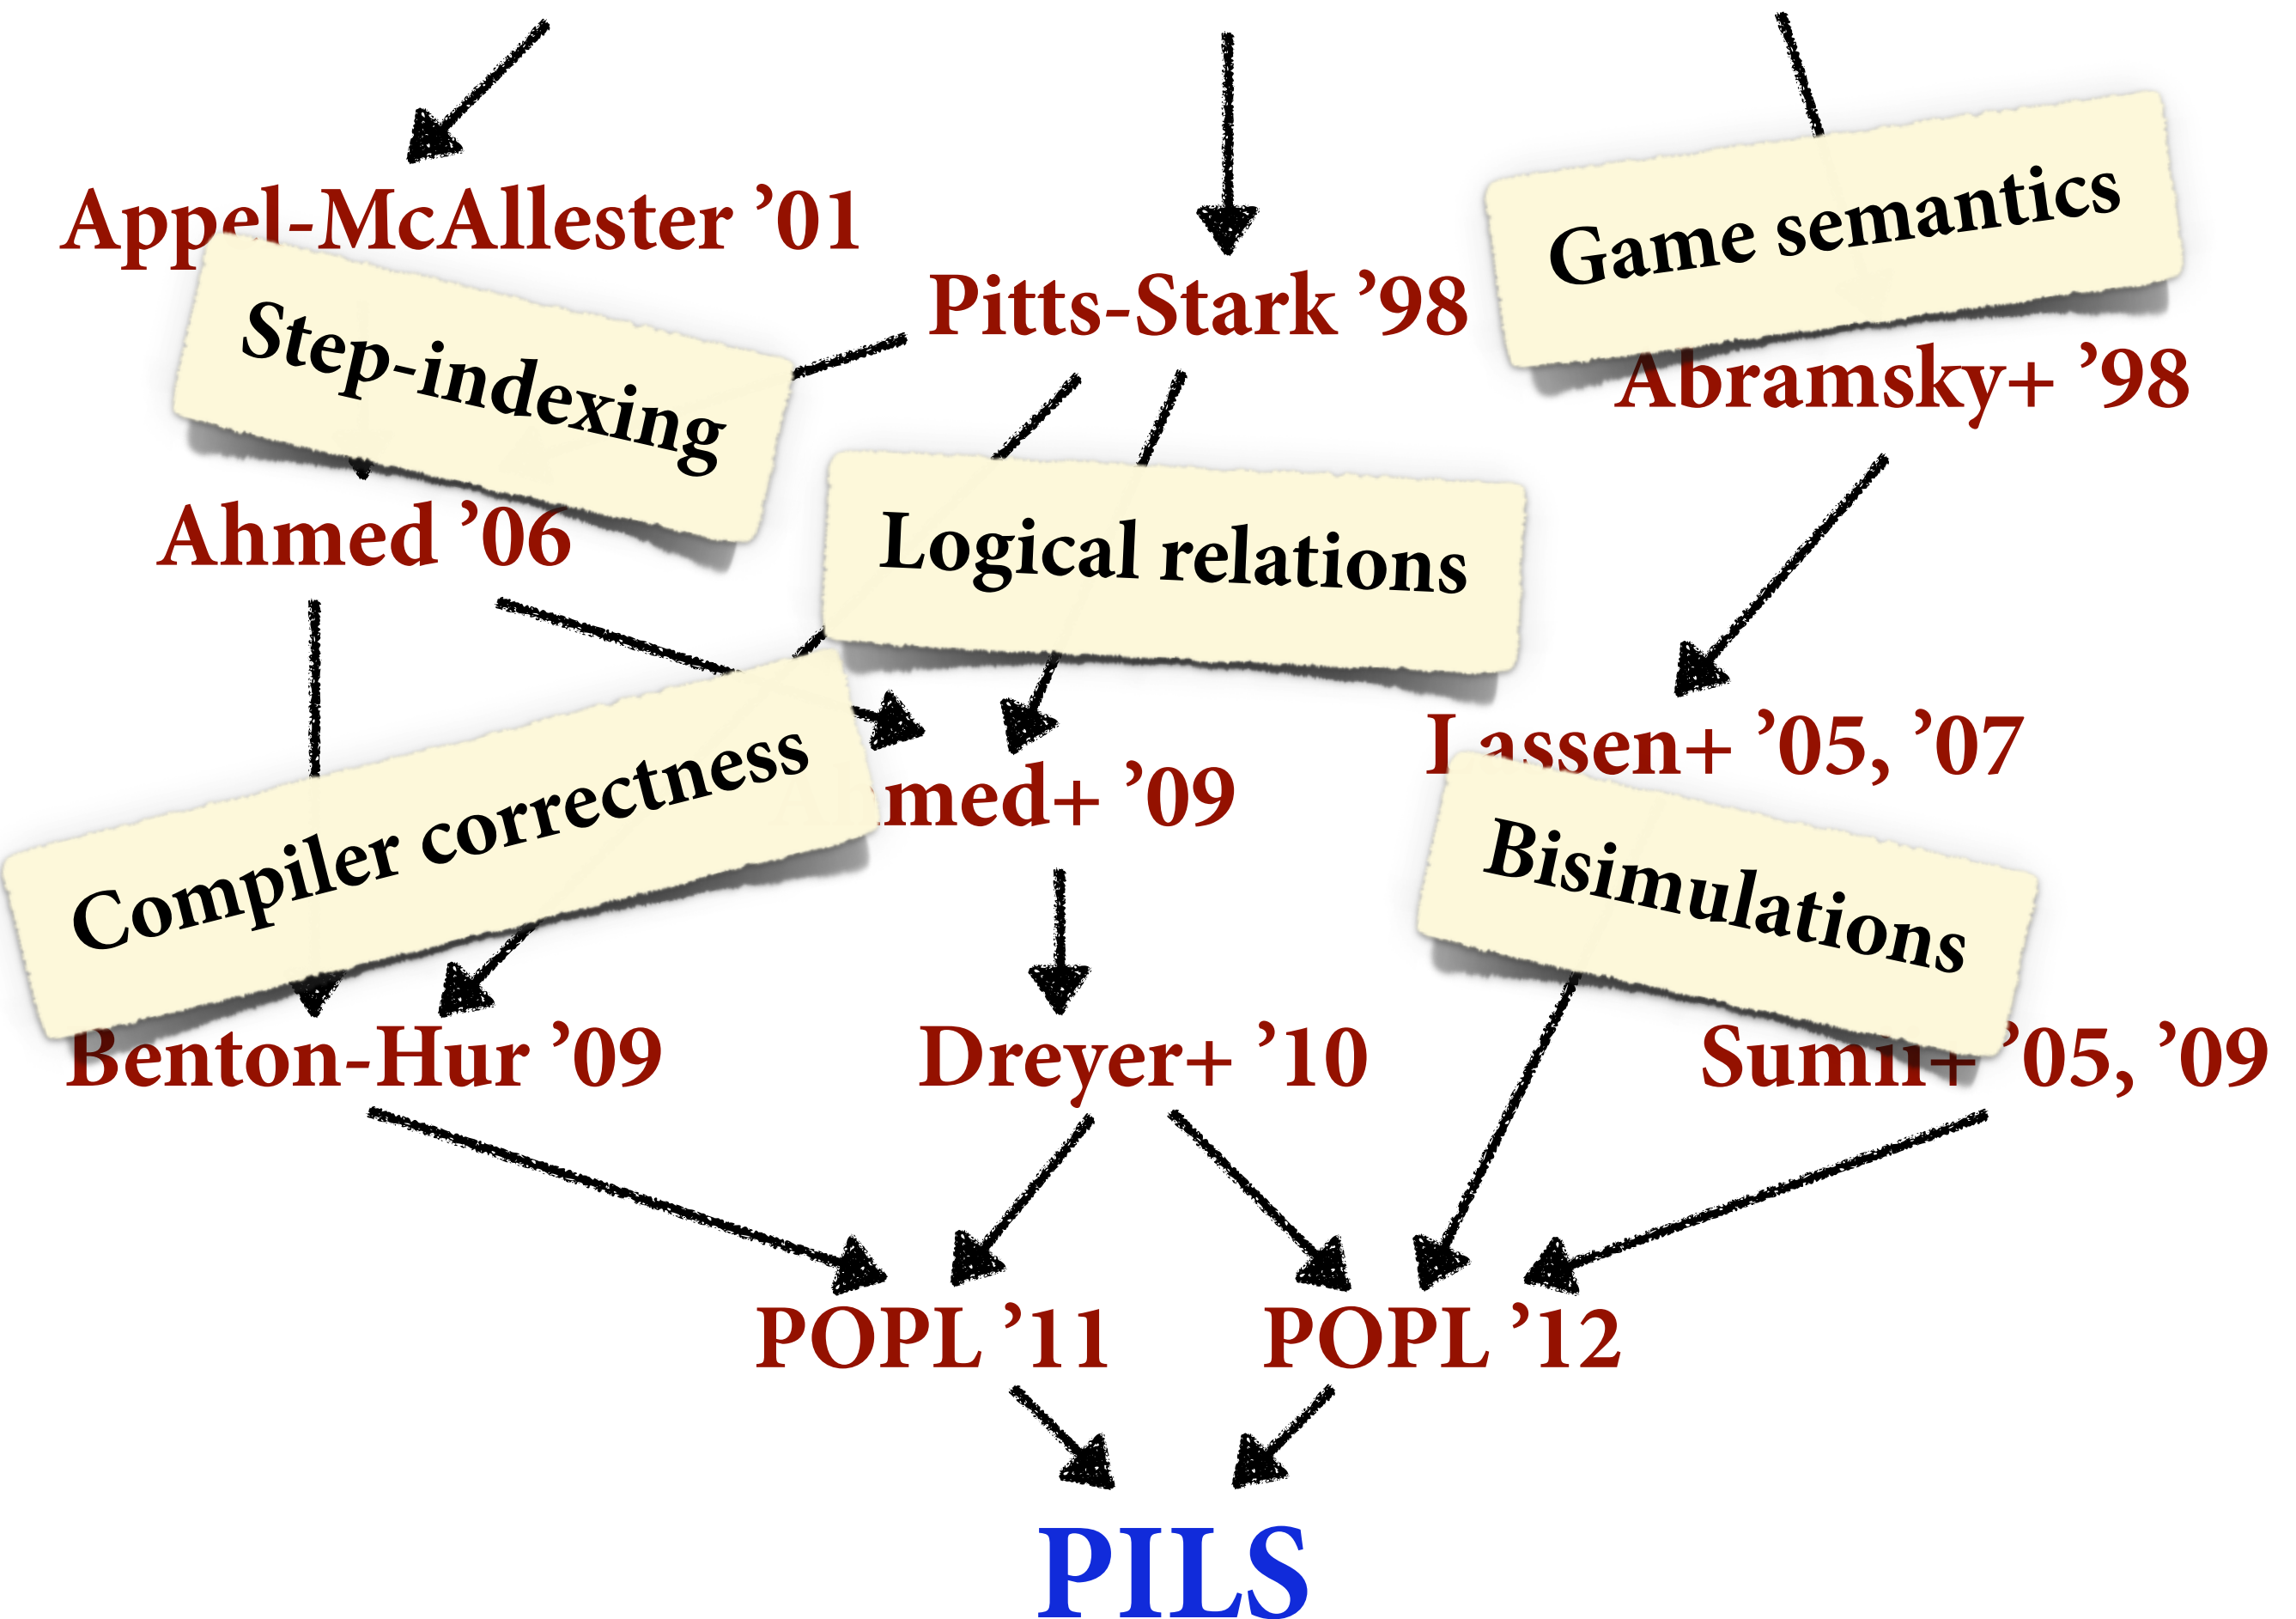
POPL '12

PILS









POPL '11

POPL '12



PILS



POPL '11

A Kripke Logical Relation Between ML and Assembly

Chung-Kil Hur* Derek Dreyer

Max Planck Institute for Software Systems (MPI-SWS)

{gil,dreyer}@mpi-sws.org

Abstract

There has recently been great progress in proving the correctness of compilers for increasingly realistic languages with increasingly realistic runtime systems. Most work on this problem has focused on proving the correctness of a particular compiler, leaving open the question of how to verify the correctness of assembly code that is hand-optimized or linked together from the output of multiple compilers. This has led Benton and other researchers to propose more abstract, compositional notions of when a low-level program correctly realizes a high-level one. However, the state of the art in so-called “compositional compiler correctness” has only considered relatively simple high-level and low-level languages.

In this paper, we propose a novel, extensional, compiler-independent notion of equivalence between high-level programs in an expressive, impure ML-like λ -calculus and low-level pro-

1. Introduction

While compiler verification is an age-old problem, there has been remarkable progress in the last several years in proving the correctness of compilers for increasingly realistic languages with increasingly realistic runtime systems. Of particular note is Leroy’s Compcert project [18], in which he used the Coq proof assistant to both program and verify a multi-pass optimizing compiler from Cminor (a C-like intermediate language) to PowerPC assembly. Dargaye [13] has adapted the Compcert framework to a compiler for a pure mini-ML language, and McCreight *et al.* [19] have extended it to support interfacing with a garbage collector. Independently, Chlipala [10, 12] has developed verified compilers for both pure and impure functional core languages, the former garbage-collected, with a focus on using custom Coq tactics to provide significant automation of verification.



POPL '12

The Marriage of Bisimulations and Kripke Logical Relations

Chung-Kil Hur Derek Dreyer Georg Neis Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

{gil,dreyer,neis,viktor}@mpi-sws.org

Abstract

There has been great progress in recent years on developing effective techniques for reasoning about program equivalence in ML-like languages—that is, languages that combine features like higher-order functions, recursive types, abstract types, and general mutable references. Two of the most prominent types of techniques to have emerged are *bisimulations* and *Kripke logical relations (KLRs)*. While both approaches are powerful, their complementary advantages have led us and other researchers to wonder whether there is an essential trade-off between them. Furthermore, both approaches seem to suffer from fundamental limitations if one is interested in scaling them to inter-language reasoning.

In this paper, we propose *relation transition systems (RTSs)*, which marry together some of the most appealing aspects of KLRs and bisimulations. In particular, RTSs show how bisimulations

purpose languages like ML that combine support for functional, *value-oriented* programming (*e.g.*, higher-order functions, polymorphism, abstract data types, recursive types) with support for imperative, *effect-oriented* programming (*e.g.*, mutable state and control effects, among other things).

Fortunately, in recent years, there has been a groundswell of interest in the problem of developing effective methods for reasoning about program equivalence in ML-like languages. A variety of promising techniques have emerged [29, 36, 19, 20, 34, 33, 23, 5, 35, 12, 25], and while some of these methods are denotational, most support direct reasoning about the operational semantics of programs. In particular, there has been a healthy rivalry between techniques based on **Kripke logical relations (KLRs)** [29, 5, 26, 13, 12, 17, 37] and **bisimulations** [36, 19, 34, 33, 23, 35].

This paper is motivated by two high-level concerns:


PILS



Putting It Together

Inter-Language

Transitive



How is a conference talk
different from a paper?

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

A paper structure that works

- **Abstract**
- **Intro**
- **Key ideas**
- **Technical meat**
- **Related work**

A paper structure that works

- **Abstract**
- **Intro**
- **Key ideas**
- **Technical meat**
- **Related work**

talk

A ~~paper~~ structure that works

- ~~Abstract~~
- Intro
- Key ideas
- ~~Technical meat~~
- ~~Related work~~

Key ideas



- Use **concrete illustrative examples** and high-level intuition.
- Do **not** show the general solution!
(People can go read your paper for that.)

talk

A ~~paper~~ structure that works

- ~~Abstract~~
- Intro
- Key ideas
- ~~Technical meat~~
- ~~Related work~~

talk

A ~~paper~~ structure that works

- **Intro** (8 minutes)
- **Key ideas** (11 minutes)

talk

A ~~paper~~ structure that works

- **Intro** (8 minutes)
- **Key ideas** (11 minutes)
- **What else is in the paper** (1 minute)

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.


Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:


- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.



Stage the motivation

- **First, get to a problem.**
 - Explain a **general** version of your problem (but not too general) **in the first 2 minutes.**
- **Then, get to the problem.**
 - Motivate and **explicitly state** your **specific** problem in the next 4 minutes.
 - Limit discussion of prior work only to what is needed to explain your problem.

WHAT DID YOU DO?!?



Tell them what you did!

- **Proudly state your contributions.**
 - After the motivation, the audience eagerly wants to hear what you did. Tell them!
- **Follow immediately with a crisp statement of your key idea(s).**
 - It will give audience a take-home message, and give focus to the rest of your talk.

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

A brief diversion into
low-level writing skills

Flow



It should be clear how each sentence and paragraph relates to **the adjacent ones**

Does this text flow?

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right.

The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right.

The game-playing technique, originally proposed by Jones et al., follows a code-based approach



What does this game-playing technique have to do with what came before?

Old to new

- Begin sentences with old info
 - Creates link to earlier text
- End sentences with new info
 - Creates link to the text that follows
 - Also places new info in position of **emphasis**



Applying old-to-new

New information

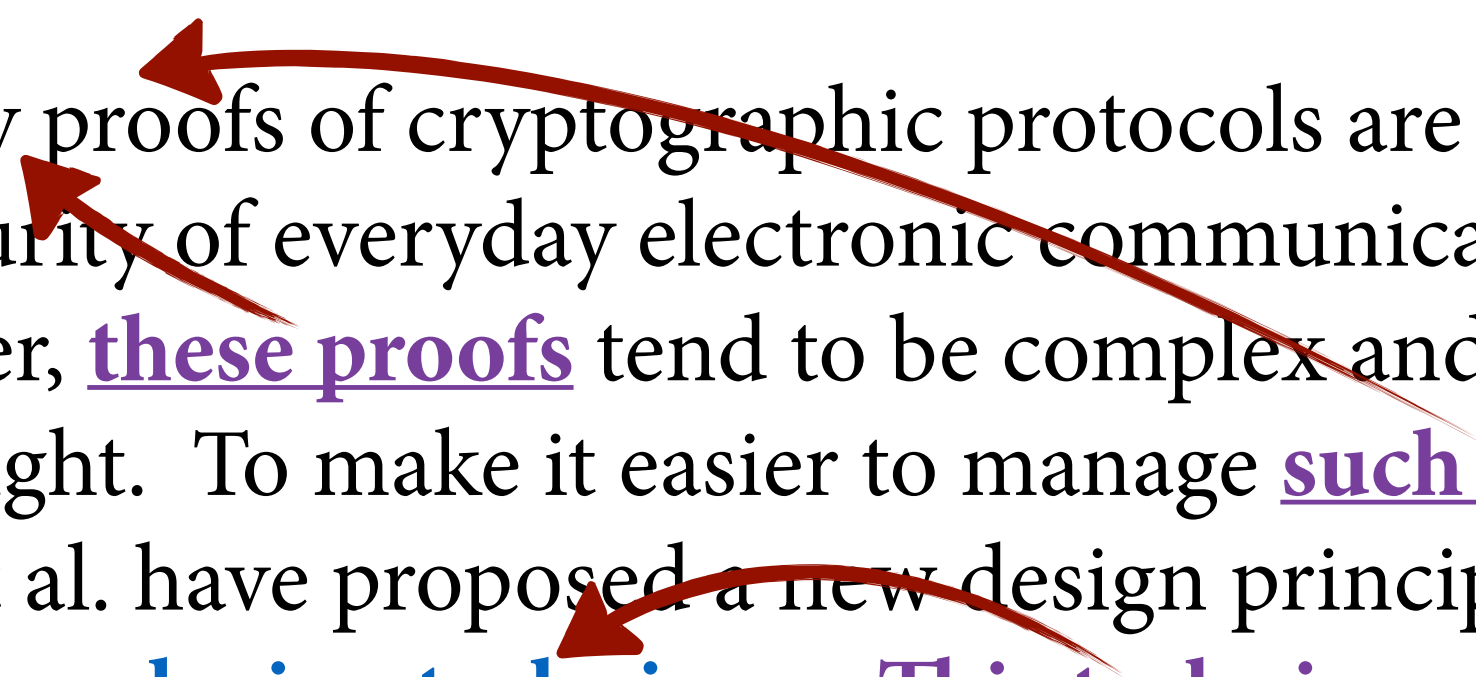
Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. **The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.** This is a general design principle for cryptographic proofs to ease their management.

Applying old-to-new

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. To make it easier to manage such proofs, Jones et al. have proposed a new design principle, called the game-playing technique. This technique follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.

Old-to-new satisfied

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. To make it easier to manage such proofs, Jones et al. have proposed a new design principle, called the **game-playing technique**. This technique follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.



The diagram consists of two red curved arrows. The first arrow starts at the underlined phrase 'these proofs' and points to the underlined phrase 'such proofs'. The second arrow starts at the underlined phrase 'This technique' and points back to the underlined phrase 'these proofs'. This visualizes a cycle of satisfaction or a transition from an old state to a new one.

But flow is not enough!

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. **Most of these large cats, however, are currently facing extinction.** A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. **A smaller cat that has been more evolutionarily successful is the house cat.** Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. **Although house cats are currently the most popular pet in the world, they are in many ways anti-social.** It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. The large cats, however, are not the only ones. A smaller, less dramatic, but equally successful species of cats are currently being studied. They are in fact, very social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

**Has great flow,
but is incoherent!**

Coherence



It should be clear how each
sentence and paragraph relates to
the big picture

One paragraph, one point

- A paragraph should have one main point, expressed in a single **point sentence**
- **Typically** the point sentence should appear **at or near the beginning of the paragraph**



Get to the
point!

No point sentence

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

Point sentence up front

There appears to be a negative correlation between the charisma of a species and its ability to survive.

Lions and tigers, for instance, are among the most majestic creatures in the animal kingdom, yet they are currently facing extinction. In contrast, the house cat is evolutionarily quite successful, even though it is mostly known for stupid pet tricks.

Flow & coherence



Create **flow** with **old to new**

Create **coherence** with
one paragraph, one point



How do flow & coherence
apply in giving talks?

Flow in talks

- **Within** a slide:
 - Script should follow “old to new”
- **Between** slides:
 - Don’t just flip to next slide and say, “So...”
 - Plan something to say **during** the transition

Flow & coherence



Create **flow** with **old to new**

Create **coherence** with
one paragraph, one point



Flow & coherence



Create flow with **old to new**

Create coherence with
~~one paragraph, one point~~
slide



Optimization & Concurrency

- Compiler performs several optimizations to generate optimized code.
 - >100 optimizations in GCC, LLVM.

Correct optimizations for sequential programs may be incorrect for shared memory concurrency.

State-of-the-Art:

- Compilers are over-conservative;
 - * optimization opportunities are lost.

or

- Buggy optimization
 - * *“Premature optimization is the root of all evil”* ~ Donald Knuth

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)
 - But just having talklets is not enough...
- **Use transitions between talklets to remind the audience of the big picture.**
 - Summarize the point of the last talklet and how it connects to the next one.

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!

On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

Conference talks

On the plus side:

- ✓ Great advertising for you and your work!


On the minus side:

- ✗ You can't say much.
- ✗ The audience may or may not care.
- ✗ Even those who care will easily get lost.
- ✗ Slides are a visual medium.

Core Idea of Rust

Unrestricted mutation and aliasing lead to:

- use-after-free errors (dangling references)
- data races
- iterator invalidation



Rust prevents all these errors using
a sophisticated “**ownership**” type system

Make the
focus obvious!




(h/t Ranjit Jhala, “How to Design Talks”)


Core Idea of Rust




Core Idea of Rust




Core Idea of Rust



Core Idea of Rust




Core Idea of Rust



Unrestricted mutation and aliasing lead to:

- use-after-free errors (dangling references)
- data races
- iterator invalidation

Core Idea of Rust



Rust prevents all these errors using a sophisticated “ownership” type system

One exception to the rule...

Talklets

Talklets

- **Break long stretches of talk into talklets.**

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)
 - But just having talklets is not enough...

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)
 - But just having talklets is not enough...
- **Use transitions between talklets to remind the audience of the big picture.**

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)
 - But just having talklets is not enough...
- **Use transitions between talklets to remind the audience of the big picture.**
 - Summarize the point of the last talklet and how it connects to the next one.

Talklets

- **Break long stretches of talk into talklets.**
 - More digestible units of story (2-4 min.)
 - But just having talklets is not enough...
- **Use transitions between talklets to remind the audience of the big picture.**
 - Summarize the point of the last talklet and how it connects to the next one.

Make the focus obvious

DO:

- Build slide visuals incrementally
- Use smooth animations to clarify transitions

DON'T:

- Reveal bullet points one at a time

Access control is inadequate, scenario 2: Facebook timeline

- Facebook introduced timeline in 2011 end
 - Chronologically order all the information on your profile
 - Make them easily searchable for other users
- Easier to search Potentially embarrassing older content
- Users were afraid of privacy violation

Access control was not changed !

Access control is inadequate, scenario 3: Spokeo

- Service aggregating information about individuals
 - Each individual information is public content
 - E.g., your Facebook profile, address
- One can infer new non public information
 - Estimating wealth using address and public property records
- Users complain of privacy violation

Access control was not changed !

Access control is inadequate: Summary

- User reaction suggests each of the cases violate privacy
- However in none of the cases access control is violated
- We propose a new model to reason about privacy

Exposure : Definition

- We define Prominence of information I at time t or $P_I(t)$
 $P_I(t) = \{U \mid U \text{ is aware of } I \text{ at time } t\}$
- Then E_I exposure of I is:

$$E_I = \lim_{t \rightarrow \infty} P_I(t)$$



Modeling user privacy using exposure

- For each content users have an expected exposure
 - How many other users are likely to access the content
- We can model privacy violation for an information as
 - Large deviation of actual exposure from expected exposure

Revisiting scenario 1: Facebook newsfeed

- Before newsfeed was introduced
 - Expected exposure: Friends who will visit user's profile
 - Actual exposure was same as expected exposure
- After newsfeed was introduced
 - Actual exposure: All friends to whom the information is pushed
 - Actual exposure is much higher than the expected exposure

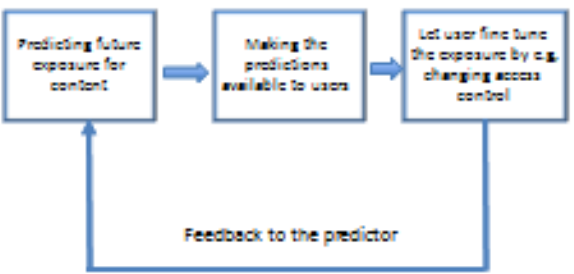
Revisiting scenario 2: Facebook timeline

- Before timeline was introduced
 - Expected exposure for older data: Friends who will scroll to find a old content
 - Actual exposure for older data was same as expected exposure
- After timeline was introduced
 - Actual exposure for older data: All friends who visit the profile
 - Actual exposure is much higher than the expected exposure

Revisiting scenario 3: Spokeo

- Before spokeo aggregated data
 - Expected exposure for new inferred data: Users who dig up each individual pieces of content form different sources
 - Actual exposure for older data was same as expected exposure
- After spokeo aggregated data
 - Actual exposure for new inferred data: All users who visit public spokeo website
 - Actual exposure is much higher than the expected exposure

Proposed model: managing privacy via exposure



Key challenge: Predicting future exposure

- Huge existing work for predicting growth in content popularity
 - Future YouTube views, Facebook likes, Retweets
 - Use machine learning, regression techniques
 - We can leverage advances in those fields to predict exposure
- OSN operators are best positioned to do the predictions
 - Empirical data on how information disseminates in their sites
 - Facebook or Youtube already provide number of likes or views

Limitations of our model

- Privacy violation by inference using available data
 - It is extremely hard to enumerate all possible inference
- Privacy violation using cross site prediction
 - Prediction across multiple systems
 - E.g., posting a picture taken from Facebook in twitter

Major Deviation from expected exposure can capture the privacy violations not covered by access control

Introduction

- Like an expanded version of the abstract
- Alternative approach (SPJ): Eliminate Context
 - Start with a concrete example, e.g. “Consider this Haskell code...”
 - If this works, it can be effective, but I find it often doesn't work
 - It assumes reader already knows context



53

A structure that works

- Abstract (1-2 paragraphs, 1000 readers)
- Intro (1-2 pages, 100 readers)
- **Key ideas** (2-3 pages, 50 readers)
- Technical meat (4-6 pages, 5 readers)
- Related work (1-2 pages, 100 readers)

54

“Key ideas” section



- Use **concrete illustrative examples** and high-level intuition
- Do **not** have to show the general solution (that's what the technical section is for)

55

Why have a “key ideas” section at all?



1. Forces you to have a “**takeaway**”
2. Many readers only care about the takeaway, not the technical details
3. For those who want the technical details, the key ideas are still useful as “scaffolding”

56


A confession



I don't always have a key ideas section.


57

Breadth-first traversal



58

Sometimes breadth-first doesn't work!
e.g., if explaining 3 & 4 requires first explaining subtree rooted at 2



○○○

59

POPL '17

A Promising Semantics for Relaxed-Memory Concurrency

Andreas Kiez, Chung-Kil Han*, Ori Lahav, Viktor Vafeiadis, Derek Dreyer

Abstract

Despite many years of research, it has proven very difficult to develop a memory model for concurrent programming languages that adequately balances the conflicting demands of programmers, compilers, and hardware. In this paper, we propose the first relaxed memory model that (1) achieves for a broad spectrum of features from the C++11 memory model (C11) to implementable, in the sense that it generally subsumes many standard compiler optimizations and mappings, as well as standard compilation techniques for C11 and PowerPC; (2) provides simple, yet robust, reasoning rules demonstrating the absence of “bad” test-and-set behavior; (3) supports “DRF” guarantees, ensuring that programmers who use software synchronization need not understand the full complexity of relaxed memory semantics; and (4) defines the semantics of new programs without relying on undefined behaviors, which is a prerequisite for applicability to type-safe languages like Java.

The key novel ideas behind our model in the context of progress, a formal progress property, are written in the future. Our starting point remains to read these ideas out of order. Crucially, no

memory shared by all threads. To simulate SC semantics on these architectures, one must therefore count expensive fence instructions to subvert the effects of the hardware. Recently, a number of concrete compiler optimizations—such as concurrent progress—are ordered (checked) by a naive SC semantics because they effectively reorder memory operations. Moreover, SC semantics is stronger (i.e., more restrictive) than necessary for many concurrent algorithms.

Since languages like Java and C++ have opted instead to provide relaxed (aka weak) memory models [C11, C13], which enable programmers to demand SC semantics when they need it, but which also support a range of cheaper memory operations that rule off strongly consistent and well-defined behavior for efficiency.

1.1. Criteria for a Programming Language Memory Model

Unfortunately, despite many years of research, it has proven very difficult to develop a memory model for concurrent programming languages that adequately balances the conflicting demands of programmers, compilers, and hardware. In particular, we would like to find a memory model that satisfies the following properties:

60

Layering the presentation

“The paper is extremely well written.”

“The presentation of the semantics is well-motivated and understandable.”

• Section 3-4: Presented other key ideas and built up to the full semantics incrementally

○○○

62

Layering the presentation

Figure 2. Operational semantics for the simplified model building only relaxed read and write actions

• What if you don't have enough space for such a layered presentation?

- Move some technical details to appendix
- Submit to a better conference (i.e. a conference with a higher page limit)

63

A structure that works

- Abstract (1-2 paragraphs, 1000 readers)
- Intro (1-2 pages, 100 readers)
- Key ideas (2-3 pages, 50 readers)
- Technical meat (4-6 pages, 5 readers)
- **Related work** (1-2 pages, 100 readers)

64

Key takeaways

- **Avoid PowerPoint-itis**

- Don't put lots of text on slides just so they are readable independently of the talk

- **Vary the look of the slides**

- Some text-only slides are fine, but if there are too many in a row, audience falls asleep

Summary of principles

- Talk \neq Paper
- Intro & key ideas are all you need
- First general problem, then specific problem
- State contributions & follow with key ideas
- Flow via old-to-new
- Coherence via one slide, one point
- Make the focus obvious
- Avoid lots of text & vary the look of slides

Summary of principles

- Talk \neq Paper
- Intro & key ideas
- First generalization
- State contributions
- Flow via old-to-new
- Coherence via one slide, one point
- Make the focus obvious
- Avoid lots of text & vary the look of slides

This is what you call
avoiding lots of text?





That's all Folks!