

Intel's Successes with Formal Methods

John Harrison

Intel Corporation

15 March 2012

Summary

- ▶ Intel's diverse verification problems

Summary

- ▶ Intel's diverse verification problems
- ▶ Verifying hardware with FEV and STE

Summary

- ▶ Intel's diverse verification problems
- ▶ Verifying hardware with FEV and STE
- ▶ Verifying protocols with model checking and SMT

Summary

- ▶ Intel's diverse verification problems
- ▶ Verifying hardware with FEV and STE
- ▶ Verifying protocols with model checking and SMT
- ▶ Verifying floating-point firmware with HOL

Summary

- ▶ Intel's diverse verification problems
- ▶ Verifying hardware with FEV and STE
- ▶ Verifying protocols with model checking and SMT
- ▶ Verifying floating-point firmware with HOL
- ▶ Perspectives and future prospects

Overview

A diversity of activities

Intel is best known as a hardware company, and hardware is still the core of the company's business. However this entails much more:

- ▶ Microcode
- ▶ Firmware
- ▶ Protocols
- ▶ Software

A diversity of activities

Intel is best known as a hardware company, and hardware is still the core of the company's business. However this entails much more:

- ▶ Microcode
- ▶ Firmware
- ▶ Protocols
- ▶ Software

If the Intel® Software and Services Group (SSG) were split off as a separate company, it would be in the top 10 software companies worldwide.

A diversity of verification problems

This gives rise to a corresponding diversity of verification problems, and of verification solutions.

- ▶ Propositional tautology/equivalence checking (FEV)
- ▶ Symbolic simulation
- ▶ Symbolic trajectory evaluation (STE)
- ▶ Temporal logic model checking
- ▶ Combined decision procedures (SMT)
- ▶ First order automated theorem proving
- ▶ Interactive theorem proving

Most of these techniques (trading automation for generality / efficiency) are in active use at Intel.

A spectrum of formal techniques

Traditionally, formal verification has been focused on complete proofs of functional correctness.

But recently there have been notable successes elsewhere for 'semi-formal' methods involving abstraction or more limited property checking.

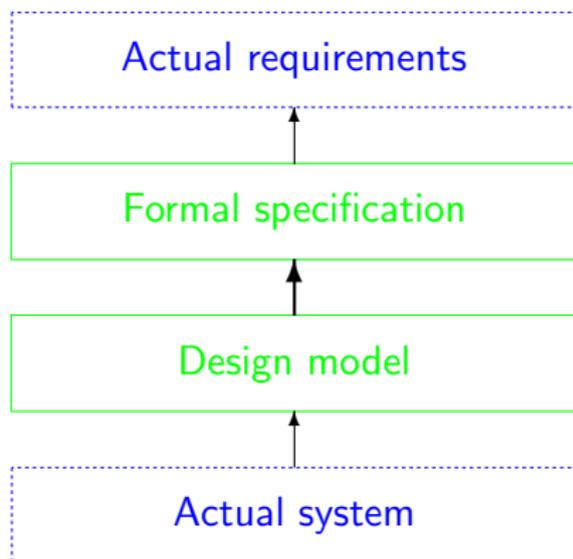
- ▶ Airbus A380 avionics
- ▶ Microsoft SLAM/SDV

One can also consider applying theorem proving technology to support testing or other traditional validation methods like path coverage.

These are all areas of interest at Intel.

Models and their validation

We have the usual concerns about validating our specs, but also need to pay attention to the correspondence between our models and physical reality.



Physical problems

Chips can suffer from physical problems, usually due to overheating or particle bombardment ('soft errors').

Physical problems

Chips can suffer from physical problems, usually due to overheating or particle bombardment ('soft errors').

- ▶ In 1978, Intel encountered problems with 'soft errors' in some of its DRAM chips.

Physical problems

Chips can suffer from physical problems, usually due to overheating or particle bombardment ('soft errors').

- ▶ In 1978, Intel encountered problems with 'soft errors' in some of its DRAM chips.
- ▶ The cause turned out to be alpha particle emission from the packaging.

Physical problems

Chips can suffer from physical problems, usually due to overheating or particle bombardment ('soft errors').

- ▶ In 1978, Intel encountered problems with 'soft errors' in some of its DRAM chips.
- ▶ The cause turned out to be alpha particle emission from the packaging.
- ▶ The factory producing the ceramic packaging was on the Green River in Colorado, downstream from the tailings of an old uranium mine.

Physical problems

Chips can suffer from physical problems, usually due to overheating or particle bombardment ('soft errors').

- ▶ In 1978, Intel encountered problems with 'soft errors' in some of its DRAM chips.
- ▶ The cause turned out to be alpha particle emission from the packaging.
- ▶ The factory producing the ceramic packaging was on the Green River in Colorado, downstream from the tailings of an old uranium mine.

However, these are rare and apparently well controlled by existing engineering best practice.

The FDIV bug

Formal methods are more useful for avoiding design errors such as the infamous FDIV bug:

The FDIV bug

Formal methods are more useful for avoiding design errors such as the infamous FDIV bug:

- ▶ Error in the floating-point division (FDIV) instruction on some early Intel® Pentium® processors

The FDIV bug

Formal methods are more useful for avoiding design errors such as the infamous FDIV bug:

- ▶ Error in the floating-point division (FDIV) instruction on some early Intel® Pentium® processors
- ▶ Very rarely encountered, but was hit by a mathematician doing research in number theory.

The FDIV bug

Formal methods are more useful for avoiding design errors such as the infamous FDIV bug:

- ▶ Error in the floating-point division (FDIV) instruction on some early Intel®Pentium® processors
- ▶ Very rarely encountered, but was hit by a mathematician doing research in number theory.
- ▶ Intel eventually set aside US \$475 million to cover the costs.

The FDIV bug

Formal methods are more useful for avoiding design errors such as the infamous FDIV bug:

- ▶ Error in the floating-point division (FDIV) instruction on some early Intel®Pentium® processors
- ▶ Very rarely encountered, but was hit by a mathematician doing research in number theory.
- ▶ Intel eventually set aside US \$475 million to cover the costs.

This did at least considerably improve investment in formal verification.

Layers of verification

If we want to verify from the level of software down to the transistors, then it's useful to identify and specify intermediate layers.

Layers of verification

If we want to verify from the level of software down to the transistors, then it's useful to identify and specify intermediate layers.

- ▶ Implement high-level floating-point algorithm assuming addition works correctly.

Layers of verification

If we want to verify from the level of software down to the transistors, then it's useful to identify and specify intermediate layers.

- ▶ Implement high-level floating-point algorithm assuming addition works correctly.
- ▶ Implement a cache coherence protocol assuming that the abstract protocol ensures coherence.

Layers of verification

If we want to verify from the level of software down to the transistors, then it's useful to identify and specify intermediate layers.

- ▶ Implement high-level floating-point algorithm assuming addition works correctly.
- ▶ Implement a cache coherence protocol assuming that the abstract protocol ensures coherence.

Many similar ideas all over computing: protocol stack, virtual machines etc.

Layers of verification

If we want to verify from the level of software down to the transistors, then it's useful to identify and specify intermediate layers.

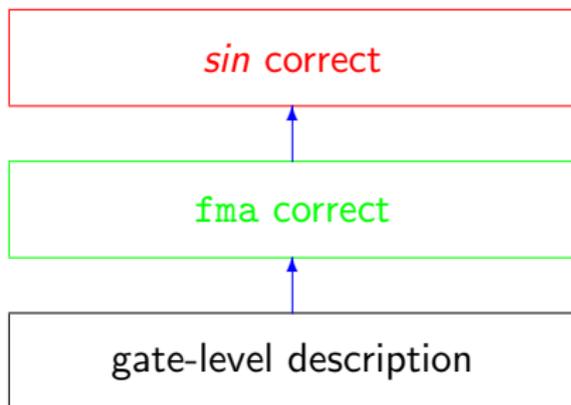
- ▶ Implement high-level floating-point algorithm assuming addition works correctly.
- ▶ Implement a cache coherence protocol assuming that the abstract protocol ensures coherence.

Many similar ideas all over computing: protocol stack, virtual machines etc.

If this clean separation starts to break down, we may face much worse verification problems. . .

How some of our verifications fit together

For example, the `fma` behavior is the *assumption* for my verification, and the *conclusion* for someone else's.



But this is not quite trivial when the verifications use different formalisms!

I: Hardware with SAT and STE

O'Leary, Zhao, Gerth, Seger, *Formally verifying IEEE compliance of floating-point hardware*, ITJ 1999.

Yang and Seger, *Introduction to Generalized Symbolic Trajectory Evaluation*, FMCAD 2002.

Schubert, *High level formal verification of next-generation microprocessors*, DAC 2003.

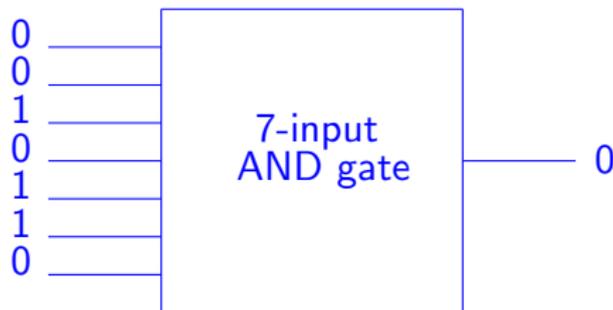
Slobodova, *Challenges for Formal Verification in Industrial Setting*, FMCAD 2007.

Kaivola et al., *Replacing Testing with Formal Verification in Intel® Core™ i7 Processor Execution Engine Validation*, CAV 2009.

Simulation

The traditional method for testing and debugging hardware designs is *simulation*.

This is just testing, done on a formal circuit model.



Feed sets of arguments in as inputs, and check whether the output is as expected.

Generalizations of simulation

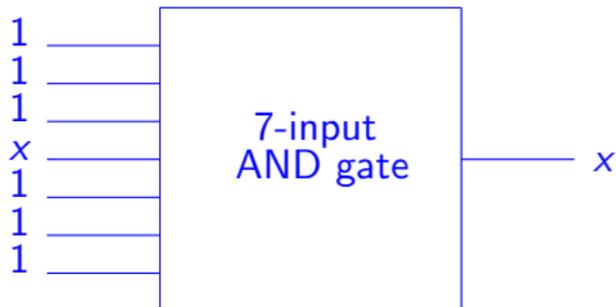
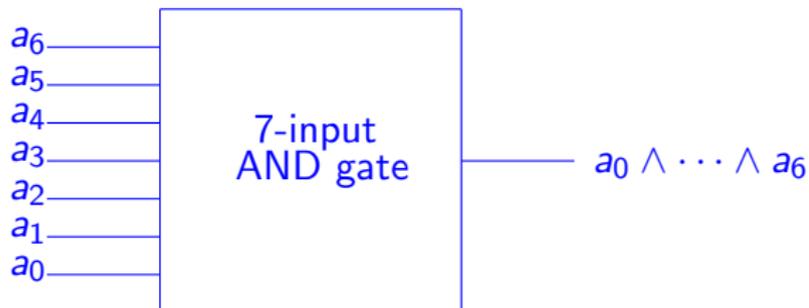
We can generalize basic simulation in two different ways:

- ▶ *Ternary* simulation, where as well as 0 and 1 we have a “don’t care” value X.
- ▶ *Symbolic* simulation, where inputs may be parametrized by Boolean variables, and outputs are functions of those variables (usually represented as BDDs).

Rather surprisingly, it’s especially useful to do both at the same time, and have ternary values parametrized by Boolean variables. This leads on to *symbolic trajectory evaluation* (STE) and its generalizations.

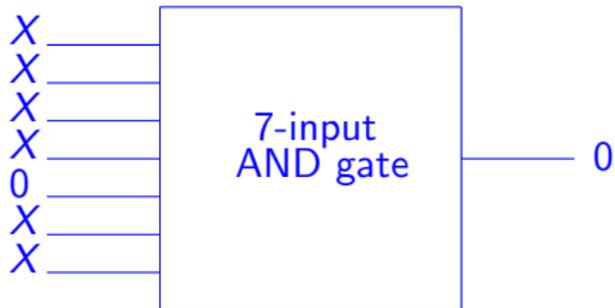
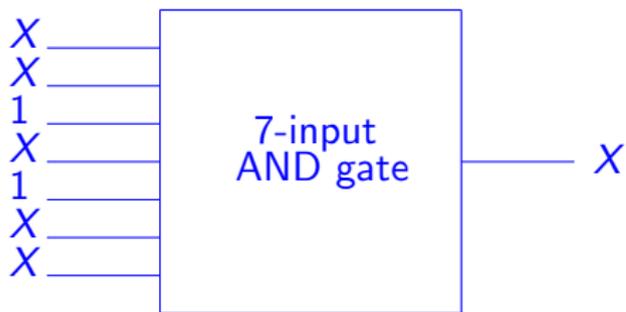
Example of symbolic simulation

We might use Boolean variables for all, or just some, inputs:



Example of ternary simulation

If some inputs are undefined, the output often is too, but not always:



Economies

Consider the 7-input AND gate. To verify it exhaustively:

- ▶ In conventional simulation, we would need **128** test cases, `0000000`, `0000001`, ..., `1111111`.
- ▶ In symbolic simulation, we only need **1** symbolic test case, `a0a1a2a3a4a5a6`, but need to manipulate expressions, not just constants.
- ▶ In ternary simulation, we need **8** test cases, `XXXXXX0`, `XXXXXX1`, ..., `0XXXXXX` and `1111111`.

If we combine symbolic and ternary simulation, we can *parametrize* the 8 test cases by just 3 Boolean variables.

This makes the manipulation of expressions much more economical.

Quaternary simulation

It's theoretically convenient to generalize ternary to *quaternary* simulation, introducing an 'overconstrained' value T .

We can think of each quaternary value as standing for a set of possible values:

$$T = \{\}$$

$$0 = \{0\}$$

$$1 = \{1\}$$

$$X = \{0, 1\}$$

This is essentially a simple case of an abstraction mapping, and we can think of the abstract values partially ordered by information.

Extended truth tables

The truth-tables for basic gates are extended:

p	q	$p \wedge q$	$p \vee q$	$p \Rightarrow q$	$p \Leftrightarrow q$
X	X	X	X	X	X
X	0	0	X	X	X
X	1	X	1	1	X
0	X	0	X	1	X
0	0	0	0	1	1
0	1	0	1	1	0
1	X	X	1	X	X
1	0	0	1	0	0
1	1	1	1	1	1

Composing gates in this simple way, we may lose information.

Symbolic trajectory evaluation

Symbolic trajectory evaluation (STE) is a further development of ternary symbolic simulation.

The user can write specifications in a restricted *temporal logic*, specifying the behavior over bounded-length *trajectories* (sequences of circuit states).

A typical specification would be: if the current state satisfies a property P , then after n time steps, the state will satisfy the property Q .

The circuit can then be checked against this specification by symbolic quaternary simulation.

STE plus theorem proving

STE (sometimes its extension GSTE) is the basic hardware verification workhorse at Intel

However, it often needs to be combined with theorem-proving for effective problem decomposition.

Intel has its own custom tool integrating lightweight theorem proving with STE, GSTE and other model checking engines.

This combination has been applied successfully to many hardware components, including floating-point units and many others.

II: Protocols with model checking and SMT

Chou, Mannava and Park: *A simple method for parameterized verification of cache coherence protocols*, FMCAD 2004.

Krstic, *Parametrized System Verification with Guard Strengthening and Parameter Abstraction*, AVIS 2005.

Talupur, Krstic, O'Leary and Tuttle, *Parametric Verification of Industrial Strength Cache Coherence Protocols*, DCC 2008.

Bingham, *Automatic non-interference lemmas for parameterized model checking*, FMCAD 2008.

Talupur and Tuttle, *Going with the Flow: Parameterized Verification Using Message Flows*, FMCAD 2008.

Parametrized systems

We mean by a *parameterized system* a transition system with N 'equivalent' replicated components, so the state space involves some Cartesian product

Parametrized systems

We mean by a *parameterized system* a transition system with N 'equivalent' replicated components, so the state space involves some Cartesian product

$$\Sigma = \Sigma_0 \times \overbrace{\Sigma_1 \times \cdots \times \Sigma_1}^{N \text{ times}}$$

Parametrized systems

We mean by a *parameterized system* a transition system with N 'equivalent' replicated components, so the state space involves some Cartesian product

$$\Sigma = \Sigma_0 \times \overbrace{\Sigma_1 \times \cdots \times \Sigma_1}^{N \text{ times}}$$

and the transition relation is symmetric between the replicated components.

Parametrized systems

We mean by a *parameterized system* a transition system with N 'equivalent' replicated components, so the state space involves some Cartesian product

$$\Sigma = \Sigma_0 \times \overbrace{\Sigma_1 \times \cdots \times \Sigma_1}^{N \text{ times}}$$

and the transition relation is symmetric between the replicated components.

Sometimes we have subtler symmetry, but we'll just consider full symmetry.

Multiprocessors with private cache

Example: multiprocessor where each processor has its own cache.
We have N cacheing agents with state space Σ_1 each, and maybe some special 'home node' with state space Σ_0 .
We can consider Σ_1 as finite with two radical but not unreasonable simplifications:

Multiprocessors with private cache

Example: multiprocessor where each processor has its own cache.

We have N cacheing agents with state space Σ_1 each, and maybe some special 'home node' with state space Σ_0 .

We can consider Σ_1 as finite with two radical but not unreasonable simplifications:

- ▶ Assume all cache lines are independent (no resource allocation conflicts)

Multiprocessors with private cache

Example: multiprocessor where each processor has its own cache.

We have N cacheing agents with state space Σ_1 each, and maybe some special 'home node' with state space Σ_0 .

We can consider Σ_1 as finite with two radical but not unreasonable simplifications:

- ▶ Assume all cache lines are independent (no resource allocation conflicts)
- ▶ Ignore actual data and consider only state of cache line (dirty, clean, whatever)

Cache coherence protocol

- ▶ A cache coherence protocol is a parametrized transition system designed to ensure that all caches have a coherent view of memory.

Cache coherence protocol

- ▶ A cache coherence protocol is a parametrized transition system designed to ensure that all caches have a coherent view of memory.
- ▶ On some simplifying assumptions, we can express this adequately using just the cache states.

Cache coherence protocol

- ▶ A cache coherence protocol is a parametrized transition system designed to ensure that all caches have a coherent view of memory.
- ▶ On some simplifying assumptions, we can express this adequately using just the cache states.
- ▶ In classic MESI protocols, each cache can be in four states: Modified, Exclusive, Shared and Invalid.

Cache coherence protocol

- ▶ A cache coherence protocol is a parametrized transition system designed to ensure that all caches have a coherent view of memory.
- ▶ On some simplifying assumptions, we can express this adequately using just the cache states.
- ▶ In classic MESI protocols, each cache can be in four states: Modified, Exclusive, Shared and Invalid.
- ▶ Coherence means:
 $\forall i. \text{Cache}(i) \in \{\text{Modified}, \text{Exclusive}\}$
 $\Rightarrow \forall j. j \neq i \Rightarrow \text{Cache}(j) = \text{Invalid}$

Why coherence isn't trivial

At an abstract level, each memory operation updates cache states appropriately, e.g.:

When writing a given line to cache i :

- ▶ *If state was Exclusive, change it to Modified.*
- ▶ *Otherwise change it to Exclusive and all $j \neq i$ to Invalid.*

Why coherence isn't trivial

At an abstract level, each memory operation updates cache states appropriately, e.g.:

When writing a given line to cache i :

- ▶ *If state was Exclusive, change it to Modified.*
- ▶ *Otherwise change it to Exclusive and all $j \neq i$ to Invalid.*

In practice, each such operation is decomposed into a complicated series of communications ('invalidate request', ...) and responses. It's very hard to be sure that with all the intricate ways these multiple transactions can interact, coherence is maintained.

Proving cache coherence

Given a protocol specification, how might we show that it ensures coherence?

Proving cache coherence

Given a protocol specification, how might we show that it ensures coherence?

- ▶ Enumeration — if the state space is finite and small enough, can analyze exhaustively (model checking).

Proving cache coherence

Given a protocol specification, how might we show that it ensures coherence?

- ▶ Enumeration — if the state space is finite and small enough, can analyze exhaustively (model checking).
- ▶ Abstraction — cleverly replace the transition system by a simpler one and infer results from that.

Proving cache coherence

Given a protocol specification, how might we show that it ensures coherence?

- ▶ Enumeration — if the state space is finite and small enough, can analyze exhaustively (model checking).
- ▶ Abstraction — cleverly replace the transition system by a simpler one and infer results from that.
- ▶ Induction — prove that invariant holds in initial states and is preserved by transitions.

Model checking

If $\Sigma_0 \times \Sigma_1^N$ is small enough, then it may be feasible to perform model checking. However:

Model checking

If $\Sigma_0 \times \Sigma_1^N$ is small enough, then it may be feasible to perform model checking. However:

- ▶ For a complex protocol (e.g. the one in Intel® QuickPath Interconnect protocol) this may be infeasible even for small N .

Model checking

If $\Sigma_0 \times \Sigma_1^N$ is small enough, then it may be feasible to perform model checking. However:

- ▶ For a complex protocol (e.g. the one in Intel® QuickPath Interconnect protocol) this may be infeasible even for small N .
- ▶ In general, a protocol is designed to work for *any* N , and we would prefer a general proof that doesn't require a specific N .

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.
- ▶ Preserved by permitted state transitions
 $P(\sigma) \wedge R(\sigma, \sigma') \Rightarrow P(\sigma')$

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.
- ▶ Preserved by permitted state transitions
 $P(\sigma) \wedge R(\sigma, \sigma') \Rightarrow P(\sigma')$

For a non-trivial protocol, that seldom happens. The solution is familiar from 'strengthening the inductive hypothesis'. We seek an I such that

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.
- ▶ Preserved by permitted state transitions
 $P(\sigma) \wedge R(\sigma, \sigma') \Rightarrow P(\sigma')$

For a non-trivial protocol, that seldom happens. The solution is familiar from 'strengthening the inductive hypothesis'. We seek an I such that

- ▶ I is an inductive invariant

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.
- ▶ Preserved by permitted state transitions
 $P(\sigma) \wedge R(\sigma, \sigma') \Rightarrow P(\sigma')$

For a non-trivial protocol, that seldom happens. The solution is familiar from 'strengthening the inductive hypothesis'. We seek an I such that

- ▶ I is an inductive invariant
- ▶ $\forall \sigma. I(\sigma) \Rightarrow P(\sigma)$

Proving coherence by induction

If we're very lucky, the property P of coherence may itself be an inductive invariant:

- ▶ Holds in the initial state $P(\sigma_0)$.
- ▶ Preserved by permitted state transitions
 $P(\sigma) \wedge R(\sigma, \sigma') \Rightarrow P(\sigma')$

For a non-trivial protocol, that seldom happens. The solution is familiar from 'strengthening the inductive hypothesis'. We seek an I such that

- ▶ I is an inductive invariant
- ▶ $\forall \sigma. I(\sigma) \Rightarrow P(\sigma)$

This may be quite difficult and need considerable skill and ingenuity.

A more frivolous example

Three bugs are crawling on the coordinate plane. They move one at a time, and each bug will only crawl in a direction parallel to the line joining the other two.

The bugs start out at $(0,0)$, $(3,0)$, and $(0,3)$.

- (a) Is it possible that after some time the first bug will end up back where it started, while the other two bugs switch places?
- (b) Can the bugs end up at $(1,2)$, $(2,5)$, and $(-2,3)$?

Finding an inductive invariant

The answer is *no* in both cases. How could we prove it?

The property of not being at a particular point is certainly *not* inductive.

So we need to find some inductive invariant that implies the property we want.

Solution to the bug puzzle

The oriented area of the triangle formed by the bugs does not change.

If the bugs are at (a_x, a_y) , (b_x, b_y) and (c_x, c_y) , the oriented area is:

$$((b_x - a_x) \cdot (c_y - a_y) - (c_x - a_x) \cdot (b_y - a_y))/2$$

In the initial configuration this is $9/2$, and in the two final configurations it's $-9/2$ and ± 5 .

So all we need to prove is that 'the oriented area is $9/2$ ' is inductive, and we're done.

This is just algebra.

Chou-Mannava-Park method

One practical approach that has been used extensively at Intel:

- ▶ Method due to Chou, Mannava and Park
- ▶ Draws inspiration from McMillan's work
- ▶ Made more systematic by Krstic
- ▶ Further generalized, extended and applied by Bingham, Talupur, Tuttle and others.

Basic idea of the method

Consider an abstraction of the system as a product of isomorphic finite-state systems parametrized by a *view*, which is a 2-element set of node indices.

Basically, for each pair of nodes $\{i, j\}$, we modify the real system by:

Basic idea of the method

Consider an abstraction of the system as a product of isomorphic finite-state systems parametrized by a *view*, which is a 2-element set of node indices.

Basically, for each pair of nodes $\{i, j\}$, we modify the real system by:

- ▶ Using as node indices the two elements i and j plus one additional node **Other**.

Basic idea of the method

Consider an abstraction of the system as a product of isomorphic finite-state systems parametrized by a *view*, which is a 2-element set of node indices.

Basically, for each pair of nodes $\{i, j\}$, we modify the real system by:

- ▶ Using as node indices the two elements i and j plus one additional node **Other**.
- ▶ Conservatively interpreting the transition relation, using **Other** in place of the 'ignored' nodes.

Basic idea of the method

Consider an abstraction of the system as a product of isomorphic finite-state systems parametrized by a *view*, which is a 2-element set of node indices.

Basically, for each pair of nodes $\{i, j\}$, we modify the real system by:

- ▶ Using as node indices the two elements i and j plus one additional node **Other**.
- ▶ Conservatively interpreting the transition relation, using **Other** in place of the 'ignored' nodes.

Too crude to deduce the desired invariant, but it is supplemented with *noninterference lemmas* in an interactive process.

Symmetric between components of the Cartesian product, so only need consider a finite-state system.

III: Floating-point firmware with HOL

Harrison, *A Machine-Checked Theory of Floating Point Arithmetic*, TPHOLs 1999.

Harrison, *Formal verification of IA-64 division algorithms*, TPHOLs 2000.

Harrison, *Formal verification of floating point trigonometric functions*, FMCAD 2000.

Harrison, *Floating-Point Verification using Theorem Proving*, SFM summer school 2006.

Our work

We have formally verified correctness of various floating-point algorithms.

- ▶ Division and square root (Marstein-style, using fused multiply-add to do Newton-Raphson or power series approximation with delicate final rounding).
- ▶ Transcendental functions like *log* and *sin* (table-driven algorithms using range reduction and a core polynomial approximations).

Proofs use the HOL Light prover

- ▶ <http://www.cl.cam.ac.uk/users/jrh/hol-light>

Our HOL Light proofs

The mathematics we formalize is mostly:

- ▶ Elementary number theory and real analysis
- ▶ Floating-point numbers, results about rounding etc.

Needs several special-purpose proof procedures, e.g.

- ▶ Verifying solution set of some quadratic congruences
- ▶ Proving primality of particular numbers
- ▶ Proving bounds on rational approximations
- ▶ Verifying errors in polynomial approximations

Example: tangent algorithm

- ▶ The input number X is first reduced to r with approximately $|r| \leq \pi/4$ such that $X = r + N\pi/2$ for some integer N . We now need to calculate $\pm \tan(r)$ or $\pm \cot(r)$ depending on N modulo 4.

Example: tangent algorithm

- ▶ The input number X is first reduced to r with approximately $|r| \leq \pi/4$ such that $X = r + N\pi/2$ for some integer N . We now need to calculate $\pm \tan(r)$ or $\pm \cot(r)$ depending on N modulo 4.
- ▶ If the reduced argument r is still not small enough, it is separated into its leading few bits B and the trailing part $x = r - B$, and the overall result computed from $\tan(x)$ and pre-stored functions of B , e.g.

$$\tan(B + x) = \tan(B) + \frac{1}{\sin(B)\cos(B)} \frac{\tan(x)}{\cot(B) - \tan(x)}$$

Example: tangent algorithm

- ▶ The input number X is first reduced to r with approximately $|r| \leq \pi/4$ such that $X = r + N\pi/2$ for some integer N . We now need to calculate $\pm \tan(r)$ or $\pm \cot(r)$ depending on N modulo 4.
- ▶ If the reduced argument r is still not small enough, it is separated into its leading few bits B and the trailing part $x = r - B$, and the overall result computed from $\tan(x)$ and pre-stored functions of B , e.g.

$$\tan(B + x) = \tan(B) + \frac{1}{\frac{\sin(B)\cos(B)}{\cot(B) - \tan(x)}} \tan(x)$$

- ▶ Now a power series approximation is used for $\tan(r)$, $\cot(r)$ or $\tan(x)$ as appropriate.

Overview of the verification

To verify this algorithm, we need to prove:

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.
- ▶ The mathematical facts used to reconstruct the result from components are applicable.

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.
- ▶ The mathematical facts used to reconstruct the result from components are applicable.
- ▶ Stored constants such as $\tan(B)$ are sufficiently accurate.

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.
- ▶ The mathematical facts used to reconstruct the result from components are applicable.
- ▶ Stored constants such as $\tan(B)$ are sufficiently accurate.
- ▶ The power series approximation does not introduce too much error in approximation.

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.
- ▶ The mathematical facts used to reconstruct the result from components are applicable.
- ▶ Stored constants such as $\tan(B)$ are sufficiently accurate.
- ▶ The power series approximation does not introduce too much error in approximation.
- ▶ The rounding errors involved in computing with floating point arithmetic are within bounds.

Overview of the verification

To verify this algorithm, we need to prove:

- ▶ The range reduction to obtain r is done accurately.
- ▶ The mathematical facts used to reconstruct the result from components are applicable.
- ▶ Stored constants such as $\tan(B)$ are sufficiently accurate.
- ▶ The power series approximation does not introduce too much error in approximation.
- ▶ The rounding errors involved in computing with floating point arithmetic are within bounds.

Most of these parts are non-trivial. Moreover, some of them require more pure mathematics than might be expected.

Why mathematics?

Controlling the error in range reduction becomes difficult when the reduced argument $X - N\pi/2$ is small.

To check that the computation is accurate enough, we need to know:

How close can a floating point number be to an integer multiple of $\pi/2$?

Even deriving the power series (for $0 < |x| < \pi$):

$$\cot(x) = 1/x - \frac{1}{3}x - \frac{1}{45}x^3 - \frac{2}{945}x^5 - \dots$$

is much harder than you might expect.

Why HOL Light?

We need a general theorem proving system with:

- ▶ High standard of logical rigor and reliability
- ▶ Ability to mix interactive and automated proof
- ▶ Programmability for domain-specific proof tasks
- ▶ A substantial library of pre-proved mathematics

Other theorem provers such as ACL2, Coq and PVS have also been used for verification in this area.

Conclusions

The value of formal verification

Formal verification has contributed in many ways, and not only the obvious ones:

- ▶ Uncovered bugs, including subtle and sometimes very serious ones
- ▶ Revealed ways that algorithms could be made more efficient
- ▶ Improved our confidence in the (original or final) product
- ▶ Led to deeper theoretical understanding

This experience seems quite common.

What's missing?

- ▶ Hardware verification proofs use STE as the workhorse, but sometimes want greater theorem-proving power than the current framework provides.
- ▶ The CMP method uses model checking with an ad hoc program for doing the abstraction and the successive refinements, not formally proved correct.
- ▶ The high-level HOL verifications assumes the correctness of the basic FP operations, but this is not the same as the low-level specs used in the hardware verification.

All in all, Intel has achieved a lot in the field of FV, but we could achieve even more with a completely seamless combination of all our favorite techniques!