

SWATI: Synthesizing Word-Lengths Automatically Using Testing and Induction

*Susmit Jha
Sanjit A. Seshia*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-7

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-7.html>

February 8, 2013



Copyright © 2013, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

SWATI: Synthesizing Wordlengths Automatically Using Testing and Induction

Susmit Jha^{1,2} and Sanjit A. Seshia²

¹ Strategic CAD Labs, Intel
susmit.jha@intel.com

² UC Berkeley
jha, sseshia@eecs.berkeley.edu

Abstract. In this paper, we present an automated technique *swati*: *Synthesizing Wordlengths Automatically Using Testing and Induction*, which uses a combination of Nelder-Mead optimization based testing, and induction from examples to automatically synthesize optimal fixedpoint implementation of numerical routines. The design of numerical software is commonly done using floating-point arithmetic in design-environments such as Matlab. However, these designs are often implemented using fixed-point arithmetic for speed and efficiency reasons especially in embedded systems. The fixed-point implementation reduces implementation cost, provides better performance, and reduces power consumption. The conversion from floating-point designs to fixed-point code is subject to two opposing constraints: (i) the word-width of fixed-point types must be minimized, and (ii) the outputs of the fixed-point program must be accurate. In this paper, we propose a new solution to this problem. Our technique takes the floating-point program, specified accuracy and an implementation cost model and provides the fixed-point program with specified accuracy and optimal implementation cost. We demonstrate the effectiveness of our approach on a set of examples from the domain of automated control, robotics and digital signal processing.

1 Introduction

Numerical software forms a critical component of embedded systems such as robotics, automated control and digital signal processing. These numerical routines have two important characteristics. First, these routines are procedures that compute some mathematical functions designed ignoring precision issues of fixed-point arithmetic. Design environments such as Simulink/Stateflow and LabVIEW allow design and simulation of numerical routines using floating-point arithmetic that closely resembles the more intuitive real arithmetic. Second, the implementation of these numerical routines run in resource-constrained environments, requiring their optimization for low resource cost and high performance. It is common for embedded platforms to have processors without floating-point units due to their added cost and performance penalty. The signal processing/control engineer must thus redesign her floating-point program to instead use *fixed-point arithmetic*. Each floating-point variable and operation in the original program is simply replaced by a corresponding fixed-point variable and operation, so the basic structure of the program does not change. The tricky part of the redesign process is to find the *optimal fixed-point types*, viz., the optimal wordlengths (bit-widths) of fixed-point variables, so that the implementation on the platform is optimal — lowest cost and highest performance — *and* the resulting fixed-point program is sufficiently

accurate. The following novel contributions are made in this paper to address this problem:

- We present a new approach for inductive synthesis of fixed-point programs from floating-point versions. The novelty stems in part from our use of optimization: we not only use optimization routines to minimize fixed-point types (bit-widths of fixed-point variables), as previous approaches have, but also show how to use an optimization oracle to systematically test the program and generate input-output examples for inductive synthesis.
- We illustrate the practical effectiveness of our technique on programs drawn from the domains of digital signal processing and control theory. For the control theory examples, we not only exhibit the synthesized fixed-point programs, but also show that these programs, when integrated in a feedback loop with the rest of the system, perform as accurately as the original floating-point versions.

2 Preliminaries

Floating-point arithmetic [8] is a system for approximately representing real numbers that supports a wide range of values. It approximates a real number using a fixed number of significant digits scaled using an exponent. The floating-point system is so called because the radix point can *float* anywhere relative to the significant digits of the number. This is in contrast to fixed-point arithmetic [25] in which there are a fixed number of digits and the radix point is also fixed. Due to this feature, a floating-point representation can represent a much wider range of values with the same number of digits. The most common floating-point representation used in computers is that defined by the IEEE 754 Standard [1]. The storage layout of the floating-point numbers consist of three basic components: the sign, the exponent, and the mantissa. The storage layout of the single-precision and double-precision floating point numbers is presented in Table 1

- The *sign bit* is 0 for a positive number and 1 for a negative number. Flipping the value of this bit flips the sign of the number.
- The *mantissa*, also known as the significand, represents the precision bits of the number. It is composed of an implicit leading bit and the fraction bits. In order to maximize the quantity of representable numbers, floating-point numbers are typically stored with the radix point after the first non-zero digit. In base 2, the only possible non-zero digit is 1. Thus, we can just assume a leading digit of 1, and don't need to represent it explicitly. As a result, the mantissa has effectively 24 bits of resolution, by way of 23 fraction bits in single-precision floating-point numbers, and 53 bits of resolution, by way of 52 fractional bits in double-precision.
- The *exponent* field needs to represent both positive and negative exponents. To do this, a bias is added to the actual exponent in order to get the stored exponent. For IEEE single-precision floats, this value is 127. Thus, an exponent of 0 means that 127 is stored in the exponent field. A stored value of 200 indicates an exponent of $(200 - 127)$, or 73. Exponents of -127 (all 0s) and $+128$ (all 1s) are reserved for special numbers. For double precision, the exponent field is 11 bits, and has a bias of 1023.

Floating-point solves a number of representation problems. Fixed-point has a fixed window of representation, which limits it from representing very large or very small numbers. Floating-point, on the other hand, employs a sort of “sliding window” of precision

	Sign	Exponent	Fraction	Bias
Single Precision	1 [31]	8 [30 – 23]	23 [22 – 00]	127
Double Precision	1 [63]	11 [62 – 52]	52 [51 – 00]	1023

Table 1: Floating-point Number Layout

appropriate to the scale of the number. The range of positive floating-point numbers can be split into normalized numbers (which preserve the full precision of the mantissa), and denormalized numbers. The denormalized numbers do not have an implicit leading bit of 1 and allow representation of really small numbers but with only a portion of the fraction’s precision. The exponent of all 0s (–127) and all 1s (128) are reserved for denormalized numbers and representing infinity respectively. A complete discussion on the semantics of floating-point operations can be found in the IEEE 754 Standard [1]. A floating-point unit (FPU) is used to carry out operations on floating-point numbers such as addition, subtraction, multiplication, division and square root. FPUs are integrated with CPUs in computers but most embedded processors do not have hardware support for floating-point operations. Emulation of floating-point operations without hardware support can be very slow. Inclusion of FPUs also increases the power consumption of the processors. This has made the use of fixed-point arithmetic very common in embedded systems. In spite of the benefits of floating-point arithmetic, embedded systems often use fixed-point arithmetic to reduce resource cost and improve performance. A fixed-point number consists of a sign mode bit, an integer part and a fractional part. We denote the fixed-point type of a variable x by $\text{fx}\tau(x)$. Formally, a fixed-point type is a triple:

$$\langle \text{Signedness}, \text{IWL}, \text{FWL} \rangle.$$

The sign mode bit `Signedness` is 0 if the data is unsigned and is 1 if the data is signed. The length of the integer part is called the integer wordlength (IWL) and the length of the fractional part is called the fractional wordlength (FWL). The fixed-point wordlength (WL) is the sum of the integer wordlength and fractional wordlength; that is, $\text{WL} = \text{IWL} + \text{FWL}$. A fixed-point number with fractional word length (FWL) is scaled by a factor of $1/2^{\text{FWL}}$. For example, a fixed-point number 01110 with 0 as `Signedness`, integer wordlength of 3 and fractional wordlength of 2 represents $14 \times 1/2^2$, that is, 3.5. Converting a fixed-point number with scaling factor R to another type with scaling factor S, requires multiplying the underlying integer by R and dividing by S; that is, multiplying by the ratio R/S. For example, converting 01110 with 0 as `Signedness`, integer wordlength of 2 and fractional wordlength of 2 into a fixed-point number with 0 as `Signedness`, integer wordlength of 2 and fractional wordlength of 3 requires multiplying with $2^3/2^2$ to obtain 011100. If the scaling factor is to be reduced, the new integer will have to be rounded. For example, converting the same fixed-point number 01110 to a fixed-point number with fractional wordlength of 0 and integer wordlength of 2 yields 011, that is, 3 which is obtained by rounding down from 3.5. The operations supported by fixed-point arithmetic are the same as those in the floating-point arithmetic standard [1] but the semantics can differ on custom hardware. For example, the rounding mode for arithmetic operations could be different, and the result could be specified to saturate or overflow/underflow in case the wordlength of a variable is not sufficient to store a computed result. One complete semantics of fixed-point operation is provided with the Fixed-point Toolbox in Matlab [2]. The range of the fixed-point number is

much smaller compared to the range of floating-point numbers for the same number of bits since the radix point is fixed and no dynamic adjustment of precision is possible. Translating a floating-point program into fixed-point program is non-trivial and requires

Number systems with WL = 32	Range	Precision
Single-precision Floating-point	$\pm 10^{-41.85}$ to $10^{38.53}$	<i>Adaptive</i>
Fixed-point type: $\langle 1, 8, 24 \rangle$	$-10^{2.11}$ to $10^{2.11}$	$10^{-7.22}$
Fixed-point type: $\langle 1, 16, 16 \rangle$	$-10^{4.52}$ to $10^{4.52}$	$10^{-4.82}$
Fixed-point type: $\langle 1, 24, 8 \rangle$	$-10^{6.93}$ to $10^{6.93}$	$10^{-2.41}$

Table 2: Range of 32 bit fixed-point and floating-point numbers

careful consideration of loss of precision and range. The integer wordlengths and fractional wordlengths of the fixed-point variables need to be carefully selected to ensure that the computation remains accurate to a specified threshold.

3 Problem Definition

We introduce a simple illustrative example to explain the problem of synthesizing an optimal fixed-point program from a floating-point program, and then present the formal problem definition.

3.1 Floating-point Implementation

Given a floating-point program, we need to synthesize fixed-point type for each floating-point variable.

Example 1: The floating-point program in this example 1 takes radius as the input, and computes the corresponding area of the circle. Notice that the fixed-point program is essentially identical to the floating-point version, except that the fixed-point types of variables `mypi`, `radius`, `t` and `area` must be identified. Recall that the fixed-point type is a triple $\langle s_j, iwl_j, fwl_j \rangle$ for j -th variable where s_j denotes the Signedness of the variable, iwl_j denotes the integer wordlength and fwl_j denotes the fraction wordlength. We use $F_{fl}(X)$ to denote the floating-point program with inputs $X = \langle x_1, x_2, \dots, x_n \rangle$. $F_{fx}(X, \mathbf{fx}\tau)$ denotes the fixed-point version of the program, where the fixed-point type of a variable $x \in X$ is $\mathbf{fx}\tau(x)$. Note that the fixed-point types in $F_{fx}(X, \mathbf{fx}\tau)$ are defined by the mapping $\mathbf{fx}\tau$.

3.2 Input Domain

The context in which a fixed-point program $F_{fx}(X, \mathbf{fx}\tau)$ is executed often provides a precondition that must be satisfied by valid inputs $\langle x_1, x_2, \dots, x_n \rangle$. This defines the

Procedure 1 Floating-point program to compute circle area

Input: radius**Output:** area

```
double mypi, radius, t, area
mypi = 3.14159265358979323846
t = radius × radius
area = mypi × t
return area
```

Procedure 2 Fixed-point program to compute circle area

Input: radius, $\langle s_j, iwl_j, fwl_j \rangle$ for $j = 1, 2, 3, 4$ **Output:** area

```
fx(s1, iwl1, fwl1) mypi
fx(s2, iwl2, fwl2) radius
fx(s3, iwl3, fwl3) t
fx(s4, iwl4, fwl4) area
mypi = 3.14159265358979323846
t = radius × radius
area = mypi × t
return area
```

input domain denoted by $Dom(X)$.

Example 2: In our example of computing area of a circle, suppose that we are only interested in the radii in the range $[0.1, 2.0)$. Then, the input domain $Dom(radius)$ is

$$radius \geq 0.1 \wedge radius < 2.0$$

3.3 Correctness Condition for Accuracy

The correctness condition specifies an error function $Err(F_{fl}(X), F_{fx}(X, \mathbf{fx}\tau))$, and a maximum error threshold $maxError$. The error function and error threshold together define a bound on the “distance” between outputs generated by the floating-point and fixed-point programs respectively. An *accurate* fixed-point program is one whose error function lies within the error threshold for all inputs in the input domain. Some common error functions are:

- Absolute difference between the floating-point function and fixed-point function: $|F_{fl}(X) - F_{fx}(X, \mathbf{fx}\tau)|$
- Relative difference between the floating-point function and fixed-point function: $\left| \frac{F_{fl}(X) - F_{fx}(X, \mathbf{fx}\tau)}{F_{fl}(X)} \right|$
- Moderated relative difference: $\left| \frac{F_{fl}(X) - F_{fx}(X, \mathbf{fx}\tau)}{F_{fl}(X) + \delta} \right|$. This approaches the relative difference for $F_{fl}(X) \gg \delta$ and approaches a weighted absolute difference for $F_{fl}(X) \ll \delta$. When $F_{fl}(X)$ can be zero for some values of X , the moderated relative difference remains bounded unlike the relative difference which becomes unbounded.

The *correctness condition for accuracy* requires that for all inputs in the provided input domain $Dom(X)$, the error function $Err(F_{fl}(X), F_{fx}(X, \mathbf{fx}\tau))$ is below the specified threshold \maxError ; i.e.,

$$\forall X \in Dom(X) . Err(F_{fl}(X), F_{fx}(X, \mathbf{fx}\tau)) \leq \maxError$$

Example 3: In our running example of computing area of a circle, the error function is chosen to be relative difference, the error threshold 0.01, and thus the correctness condition is $\forall radius, s.t. radius \geq 0.1 \wedge radius < 2.0$

$$\frac{F_{fl}(\text{radius}) - F_{fx}(\text{radius}, \mathbf{fx}\tau)}{F_{fl}(\text{radius})} \leq 0.01$$

3.4 Implementation Cost Model

The *cost model* of the fixed-point program is a function mapping fixed-point types to a real number. For a given fixed-point program $F_{fx}(X, \mathbf{fx}\tau)$, let $T = \{t_1, t_2, \dots, t_k\}$ be the set of fixed-point program variables with corresponding types $\{\mathbf{fx}\tau(t_1), \mathbf{fx}\tau(t_2), \dots, \mathbf{fx}\tau(t_k)\}$. Then the cost model (or simply *cost*) of F_{fx} is a function

$$\text{cost} : (\mathbf{fx}\tau(t_1), \mathbf{fx}\tau(t_2), \dots, \mathbf{fx}\tau(t_k)) \rightarrow \mathbb{R}$$

In practice, *cost* is often just a function of the total wordlengths ($WL = IWL + FWL$) of the variables. It can incorporate hardware implementation metrics such as area, power and delay. A number of cost models are available in the literature [17, 18, 5, 6], and all of these can be used in our approach.

Example 4: The cost model proposed by Constantinides et al [6] for the running example yields the following cost function. We use this cost model in all our examples.

$$\begin{aligned} \text{cost}(\mathbf{fx}\tau(\text{mypi}), \mathbf{fx}\tau(\text{radius}), \mathbf{fx}\tau(\text{t}), \mathbf{fx}\tau(\text{area})) = \\ \text{cdelay}(WL(\text{mypi})) + \text{cmul}(WL(\text{radius}), WL(\text{radius}), WL(\text{t})) \\ + \text{cmul}(WL(\text{mypi}), WL(\text{t}), WL(\text{area})), \text{ where} \\ \text{cdelay}(l) = l + 1 \text{ and } \text{cmul}(l_1, l_2, l) = 0.6 \times (l_1 + 1) * l_2 - 0.85 * (l_1 + l_2 - l) \end{aligned}$$

The area of a multiplier grows almost linearly with both the coefficients and the data wordlength. The first term in the Constantinides model represents this cost. The second term represents the area cost of computational elements required only for carry propagation. The coefficients 0.6 and 0.85 were obtained through least-squared fitting to area of several hundred multipliers of different coefficient value and width [6].

3.5 Problem Definition

Definition 1 (Optimal Fixed-point Types Synthesis). *The optimal fixed-point types synthesis problem is as follows. Given a floating-point program $F_{fx}(X, \mathbf{fx}\tau(T))$ with variables T , an input domain $Dom(X)$, a correctness condition $Err(F_{fl}(X), F_{fx}(X, \mathbf{fx}\tau(T))) \leq \maxError$, and a cost model $\text{cost}(\mathbf{fx}\tau(t_1), \mathbf{fx}\tau(t_2), \dots, \mathbf{fx}\tau(t_k))$, the optimal fixed-point types synthesis problem is to discover fixed-point types*

$$\mathbf{fx}\tau^*(T) = \{\mathbf{fx}\tau^*(t_1), \mathbf{fx}\tau^*(t_2), \dots, \mathbf{fx}\tau^*(t_k)\}$$

such that the fixed-point program $F_{fl}(X)$ with the above types for fixed-point variables satisfies the correctness condition for accuracy, that is,

$$(a) \forall X \in Dom(X) . Err(F_{fl}(X), F_{fx}(X, \mathbf{fx}\tau^*(T))) \leq \text{maxError}$$

and has minimal cost with respect to the given cost function among all fixed-point types that satisfy condition (a), that is,

$$(b) \mathbf{fx}\tau^* = \underset{\mathbf{fx}\tau \text{ satisfies (a)}}{\operatorname{argmin}} \operatorname{cost}(\mathbf{fx}\tau(T))$$

Our goal is to automated this search for optimal fixed-point types. We illustrate this problem using the running example below.

Example 5: In our running example of computing the *area of a circle*, we need to discover $\mathbf{fx}\tau^*(\text{mypi})$, $\mathbf{fx}\tau^*(\text{radius})$, $\mathbf{fx}\tau^*(\text{t})$ and $\mathbf{fx}\tau^*(\text{area})$ such that

(a) the fixed-point program with the given fixed-point types satisfies the correctness condition; that is, $\forall \text{radius}, s.t., \text{radius} \geq 0.1 \wedge \text{radius} < 2.0$

$$\frac{F_{fl}(\text{radius}) - F_{fx}(\text{radius}, \mathbf{fx}\tau^*)}{F_{fl}(\text{radius})} \leq 0.01$$

(b) and the cost is minimized; that is,

$$\mathbf{fx}\tau^* = \underset{\mathbf{fx}\tau \text{ satisfies (a)}}{\operatorname{argmin}} \operatorname{cost}(\mathbf{fx}\tau(\text{mypi}, \text{radius}, \text{t}, \text{area}))$$

We use this example to illustrate the trade-off between cost and error and how a human might use trial and error to discover the correct wordlengths. We vary the wordlength of the variables. The integer wordlength is selected to avoid overflow and the remaining bits are used for fractional wordlength.

Case 1 (Figure 1): $WL = 8$ for all variables. $\mathbf{fx}\tau(\text{mypi}) = \langle 0, 2, 6 \rangle$, $\mathbf{fx}\tau(\text{radius}) = \langle 0, 1, 7 \rangle$, $\mathbf{fx}\tau(\text{t}) = \langle 0, 2, 6 \rangle$, $\mathbf{fx}\tau(\text{area}) = \langle 0, 4, 4 \rangle$. Cost is 81.80.

Case 2 (Figure 2): $WL = 12$ for all variables. $\mathbf{fx}\tau(\text{mypi}) = \langle 0, 2, 10 \rangle$, $\mathbf{fx}\tau(\text{radius}) = \langle 0, 1, 11 \rangle$, $\mathbf{fx}\tau(\text{t}) = \langle 0, 2, 10 \rangle$, $\mathbf{fx}\tau(\text{area}) = \langle 0, 4, 8 \rangle$. Cost is 179.80.

Case 3 (Figure 3): $WL = 16$ for all variables. $\mathbf{fx}\tau(\text{mypi}) = \langle 0, 2, 14 \rangle$, $\mathbf{fx}\tau(\text{radius}) = \langle 0, 1, 15 \rangle$, $\mathbf{fx}\tau(\text{t}) = \langle 0, 2, 14 \rangle$, $\mathbf{fx}\tau(\text{area}) = \langle 0, 4, 12 \rangle$. Cost is 316.20.

As we will show in the next section, our approach computes fixed-point types that meet the accuracy threshold and yield a cost of only 104.65, which, while being less than the cost in Case 1, satisfies the correctness criterion like Case 2. In the following section, we discuss our automated approach to solve this problem.

4 Our Approach

A central idea behind our approach, `swati` is to identify a small set of *interesting* inputs $S(X)$ using testing from the input domain $Dom(X)$ such that the optimal implementation found using induction that satisfies the correctness condition for the inputs in $S(X)$ will be optimal and correct for all inputs in the given input domain $Dom(X)$.

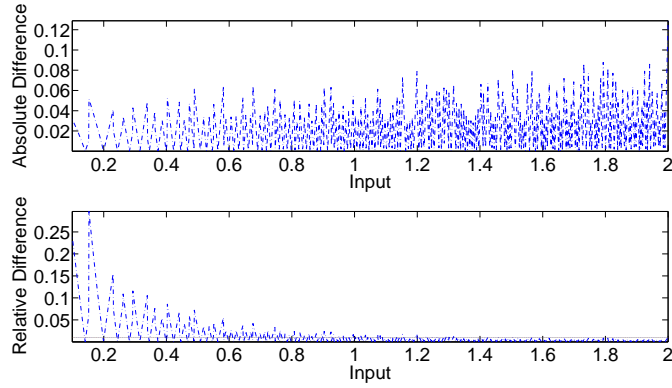


Fig. 1: Error for $WL = 8$. Error threshold at 0.01 is violated.




Fig. 2: $WL = 12$. Error threshold at 0.01 is violated.

The top-level synthesis algorithm is presented in Procedure 3. WL_{max} is an upper bound on wordlengths beyond which it is non-optimal to use the fixed-point version. The algorithm starts with a randomly selected set of examples S^0 from the given input domain. Then, a fixed-point implementation that satisfies the accuracy condition for each of these inputs and is of minimal cost is synthesized using the routine `optInduce`. If no such implementation is found, the algorithm reports `INFEASIBLE`. Otherwise, the testing routine `testErr` checks whether the implementation fails the correctness condition for any input. If so, a set of inputs Bad^i on which the implementation violates the correctness condition are added to the set S^i used for synthesis, and the process is repeated. If the correctness condition is satisfied, the resulting fixed-point types are output. In the rest of this section, we describe the main components of our approach in detail, including the theoretical result.




Fig. 3: WL = 16. Error threshold at 0.01 is not violated.

4.1 Synthesizing Optimal Types for a Finite Input Set

The `optInduce` function (see Procedure 4) is used to obtain optimum fixed-point types such that the fixed-point program with these types satisfies the correctness condition for a finite input set S and has minimal cost. First, the floating-point program F_{fl} is executed for all the inputs in the sample S and the range of each variable t_i as well as its Signedness is recorded by the functions `getRange` and `isSigned` respectively. Then, the integer wordlength IWL sufficient to represent the computed range is assigned to each variable t_i and the Signedness is 1 if the variable takes both positive and negative values, and 0 otherwise. If the fixed-point program with maximum wordlengths WL_{max} fails the correctness condition, we conclude that the synthesis is not feasible and return \perp . If not, we search for the wordlength with minimum cost satisfying the correctness condition using our optimization oracle \mathcal{O}_S . The result is used to compute the fractional wordlengths, and the resulting fixed-point types are returned.

More precisely, \mathcal{O}_S solves the following optimization problem over $\mathbf{fx}\tau$:
Minimize $\text{cost}(\mathbf{fx}\tau)$ s.t.

$$\bigwedge_{x \in S} \text{Err}(F_{fx}(x, \mathbf{fx}\tau), F_{fl}(x)) \leq \text{maxError} \quad (1)$$

Let us reflect on the nature of the above optimization problem. The overall synthesis algorithm might make several calls to \mathcal{O}_S for solving the optimization problem for different sets of inputs and hence, \mathcal{O}_S must be a fast procedure. But it is a discrete optimization problem with a non-convex constraint space, a problem class that is known to be computationally hard [7]. This rules out any computationally efficient algorithm to implement \mathcal{O}_S without sacrificing correctness guarantees. Since the space of possible types grows exponentially with the number of variables, brute-force search techniques will not scale beyond a few variables. Satisfiability solvers can also not be directly exploited to search for optimal wordlengths since the existential quantification is over the types and not the variables. The arithmetic operators have different semantics when operating on operands with different types and hence, the only way to encode this search problem as a satisfiability problem is to case-split *exhaustively* on all possible types

Procedure 3 Overall Synthesis Algorithm: `swat i`

Input: Floating-point program F_{fp} ,
Fixed-point program F_{fx} with fixed-point variables T ,
Domain of inputs Dom , Error function Err ,
maximum error threshold \maxError , Cost Model $cost$,
maximum wordlengths WL_{max}
Output: Fixed-point type $\mathbf{fx}\tau$ for variables T
or INFEASIBLE
 $S^0 = \text{random sample from } Dom, Bad^0 = S^0, i = 0$
while $Bad^i \neq \emptyset$ **do**
 $i = i + 1$
 $S^i = S^{i-1} \cup Bad^{i-1}$
 $\mathbf{fx}\tau^i = \text{optInduce}(F_{fp}, F_{fx}, Dom, Err, \maxError,$
 $cost, WL_{max}, S^i)$
 if $\mathbf{fx}\tau^i = \perp$ **then**
 return INFEASIBLE
 end if
 $Bad^i = \text{testErr}(F_{fp}, F_{fx}, \mathbf{fx}\tau^i, Dom, Err, \maxError)$
end while
return $\mathbf{fx}\tau^* = \mathbf{fx}\tau^i$

(word-lengths), where each case encodes the fixed-point program with one possible type. The number of such cases is exponential in the number of the variables in the program under synthesis and hence, SAT problems will be themselves exponentially large in size. Further, one would need to invoke SAT solvers multiple times in order to optimize the cost function. Thus, satisfiability solving would be a wrong choice to address this problem. Further, the space of possible types is also not totally ordered and hence, binary search like techniques would also not work. For a binary search like technique to work, we will need to define a domination ordering over the types which has three properties. Firstly, it is a total ordering relation. Secondly, if a particular type assignment satisfies the correctness condition for all inputs then all dominating types satisfy the correctness condition for all inputs. Thirdly, the cost function is monotonic with respect to the domination ordering relation. In general, this may not be feasible for any given floating-point program and cost function. Hence, we implement \mathcal{O}_S using a greedy procedure `getMinCostWL` presented in Procedure 5.

4.2 Verifying a Candidate Fixed-Point Program

In order to verify that the fixed-point program $F_{fx}(X, \mathbf{fx}\tau)$ satisfies the correctness condition, we need to check if the following logical formula is satisfiable.

$$\exists X \in Dom(X) \quad Err(F_{fx}(X, \mathbf{fx}\tau), F_{fp}(X)) > \maxError \quad (2)$$

If the formula is unsatisfiable, there is no input on which the fixed-point program violates the correctness condition.

For arbitrary (possibly non-linear) floating-point and fixed-point arithmetic operations, it is extremely difficult to solve such a problem in practice with current constraint solvers. Instead, we use a novel optimization-based approach to verify the candidate fixed-point program. The intuition behind using an optimization-based approach is that

Procedure 4 Optimal Fixed-Point Types Synthesis: optInduce

Input: Floating-point program F_{fp} ,
Fixed-point program F_{fx} with fixed-point variables T , Domain of inputs Dom , Error function Err ,
maximum error threshold $maxError$, Cost Model $cost$,
max wordlengths WL_{max} , Input S

Output: Optimal wordlengths WL for inputs S or \perp

```
for all fixed-point variable  $t_i$  in  $F_{fx}$  do
   $IWL(t_i) = \lceil \log(\text{getRange}(t_i, F_{fl}, S) + 1) \rceil$ 
   $Signedness(t_i) = \text{isSigned}(t_i, F_{fl}, S)$ 
end for
if  $WL_{max} < IWL$  then
  return  $\perp$ 
end if
 $fx\tau = \langle Signedness, IWL, WL_{max} - IWL \rangle$ 
if  $Err(F_{fp}(x), F_{fx}(x, fx\tau)) > maxError$  then
  return  $\perp$ 
end if
 $WL = \text{getMinCostWL}(F_{fp}, F_{fx}, Dom, Err, maxError,$ 
                      $cost, WL_{max}, S^i, IWL, Signedness)$ 
return  $fx\tau = \langle Signedness, IWL, WL - IWL \rangle$ 
```

Procedure 5 getMinCostWL

Input: Floating-point program F_{fp} ,
Fixed-point program F_{fx} with fixed-point variables T ,
Domain of inputs Dom , Error function Err ,
maximum error threshold $maxErr$, Cost Model $cost$,
max wordlengths WL_{max} , Input S

Output: Optimal wordlengths WL

```
 $valcandWL = \{WL_{max}\}$ 
while  $valcandWL$  is not empty do
   $WL = \underset{vcWL \in valcandWL}{\text{argmin}} cost(vcWL)$ 
   $fx\tau = \langle Signedness, IWL, WL - IWL \rangle$ 
   $candWL = \emptyset, valcandWL = \emptyset$ 
  for all fixed-point variable  $t_i$  in  $F_{fx}$  do
     $WL^{i-}(j) = WL(j) \forall j \neq i, WL^{i-}(i) = WL(i) - 1$ 
     $WL^{i+}(j) = WL(j) \forall j \neq i, WL^{i+}(i) = WL(i) + 1$ 
     $candWL = candWL \cup \{WL^{i-}, WL^{i+}\}$ 
  end for
  for all  $cand$  in  $candWL$  do
     $candfx\tau = \langle Signedness, IWL, candWL - IWL \rangle$ 
    if  $Err(F_{fp}(x), F_{fx}(x, cand)) \leq maxErr \forall x \in S$ 
    and  $cost(candfx\tau) < cost(fx\tau)$  then
       $valcandWL = valcandWL \cup \{cand\}$ 
    end if
  end for
end while
return  $fx\tau$ 
```

the error function is continuous in the inputs or with very few discontinuities [19, 4], and hence, optimization routines can easily find inputs which maximize error function by starting from some random input and gradually adjusting the output to increase the value of the error function. The optimization oracle \mathcal{O}_V is used to maximize the error function $Err(F_{fx}(X, \mathbf{fx}\tau), F_{fp}(X))$ over the domain $Dom(X)$. If there is no input $X \in Dom(X)$ for which the error function exceeds `maxError`, the fixed-point program is correct and we terminate. Otherwise, we obtain an example input on which the fixed-point program violates the correctness condition. Multiple inputs can also be generated where they exist.

In practice, with the current state-of-the-art optimization routines, it is difficult to implement \mathcal{O}_V to find a global optimum. Instead, we use a numerical optimization routine based on the Nelder-Mead method [21] which can handle arbitrary non-linear functions and generates local optima. Procedure 6 defines `testErr` which invokes the Nelder-Mead routine (indicated by “`argmaxlocal`”). This routine requires one to supply a starting value of X , which we generate randomly. To find multiple inputs, we invoke the routine from different random initial points and record all example inputs on which the fixed-point program violates the correctness condition. Since a global optimum is not guaranteed, we repeat this search `maxAttempts` times before declaring that the fixed-point program is correct.

Procedure 6 Verification Routine `testErr`

Input: Floating-point program F_{fp} ,
Fixed-point program F_{fx} , Fixed-point type $\mathbf{fx}\tau$,
Domain of inputs Dom , Error function Err ,
maximum error threshold `maxError`

Output: Inputs Bad on which F_{fx} violates correctness condition
 $Bad = \emptyset$

while $i \leq \text{maxAttempts}$ **do**
 $i = i + 1$, $X_0 = \text{random sample from } Dom$
 $X_{cand} = \text{argmaxlocal}(Err(F_{fp}(X), F_{fx}(X, \mathbf{fx}\tau)), X_0)$
if $Err(F_{fp}(X_{cand}), F_{fx}(X_{cand}, \mathbf{fx}\tau)) > \text{maxError}$ and $X \in Dom$ **then**
 $Bad = Bad \cup \{X\}$
end if
end while

4.3 Illustration on Running Example

The initial sample S^0 is of size 10. `maxAttempts` in the verification routine was also set to 10. The number of Our algorithm took 4 iterations. We record the intermediate implementations produced by *synthMinCost* in each of the last 3 iterations taken by our algorithm. The initial selected random sample is of size 10 and the `MAXATTEMPTS` was set to 10. The number of samples used in each subsequent step of iteration after adding examples discovered by `testErr` procedure was 18, 22 and 34. Figure 4, 5 and 6 illustrates the error using the intermediate implementations and the final implementation produced by our approach. The wordlengths in Figure 4 are *mypi*(2, 3), *radius*(1, 8), *t*(2, 11), *area*(4, 11), those in Figure 5




Fig. 4: Iteration 2




Fig. 5: Iteration 3

are $mypi(2, 5)$, $radius(1, 8)$, $t(2, 11)$, $area(4, 12)$ and in Figure 6 are $mypi(2, 3)$, $radius(1, 9)$, $t(2, 11)$, $area(4, 10)$. We observe that the number of inputs violating correctness constraint reduces after each iteration. This illustrates how our algorithm works by identifying a few representative inputs violating correctness condition in each iteration and adding that to the set of examples for which we synthesize the least cost implementation.

4.4 Theoretical Results

The following theorem summarizes the correctness and optimality guarantees of our approach.

Theorem 1. *The synthesis procedure presented in Procedure 3 is guaranteed to synthesize the fixed-point program which is of minimal cost and satisfies the correctness condition for accuracy if optimization oracles \mathcal{O}_S and \mathcal{O}_V find globally-optimal solutions (when they exist).*




Fig. 6: Iteration 4 (last)

Proof. We first prove correctness and then optimality of the obtained solution. Consider Equation 2:

$$\exists X \in \text{Dom}(X) \quad \text{Err}(F_{fx}(X, \mathbf{fx}\tau), F_{fp}(X)) > \text{maxError}$$

If \mathcal{O}_V finds globally-optimal solutions, it will find the $X \in \text{Dom}(X)$ that maximizes $\text{Err}(F_{fx}(X, \mathbf{fx}\tau), F_{fp}(X))$, and hence determine where or not Equation 2 is satisfiable. Thus, the correctness condition for accuracy is satisfied.

Next, let $\mathbf{fx}\tau^*$ ($\neq \perp$) be the fixed-point type returned by Procedure 3. Let us assume that there exists a fixed-point type $\mathbf{fx}\tau'$ with a cost lower than $\mathbf{fx}\tau^*$ which also satisfies the correctness condition:

$$\begin{aligned} & \forall X \in \text{Dom}(X) \quad \text{Err}(F_{fx}(X, \mathbf{fx}\tau'), F_{fp}(X)) \leq \text{maxError} \\ \wedge & \quad \text{cost}(\mathbf{fx}\tau') < \text{cost}(\mathbf{fx}\tau^*) \end{aligned}$$

Hence, for any $D \subseteq \text{Dom}(X)$,

$$\begin{aligned} & \forall X \in D \quad \text{Err}(F_{fx}(X, \mathbf{fx}\tau'), F_{fp}(X)) \leq \text{maxError} \\ \wedge & \quad \text{cost}(\mathbf{fx}\tau') < \text{cost}(\mathbf{fx}\tau^*) \end{aligned}$$

But $\mathbf{fx}\tau^*$ is the solution generated by applying \mathcal{O}_S to the optimization problem of Equation 1:

$$\begin{aligned} & \text{Minimize } \text{cost}(\mathbf{fx}\tau) \text{ s.t.} \\ & \bigwedge_{x \in S} \text{Err}(F_{fx}(x, \mathbf{fx}\tau), F_{fl}(x)) \leq \text{maxError} \end{aligned} \quad (3)$$

Since \mathcal{O}_S is guaranteed to generate globally-optimal solutions, setting $D = S$, we obtain a contradiction. Hence, there exists no fixed-point type $\mathbf{fx}\tau'$ with a cost lower than $\mathbf{fx}\tau^*$ which also satisfies the correctness condition. Hence, $\mathbf{fx}\tau^*$ is the optimal correct solution.

As noted earlier, it is difficult to implement ideal \mathcal{O}_S and \mathcal{O}_V (that find global optima) with current SAT and optimization methods for arbitrary floating-point programs. Nonetheless, our experience with heuristic methods that find local optima has been very good. Also, improvements in optimization/SAT methods can directly be leveraged with our inductive synthesis approach. In contrast, the current techniques for synthesizing fixed-point versions of floating-point programs perform heuristic optimization over a randomly selected set of inputs (see Sec. 6 for a detailed discussion). Such techniques do not provide any correctness guarantees and the number of inputs needed could be much larger. Our approach systematically discovers a small number of example inputs such that the optimal fixed-point program for this set yields that for the entire input domain.

5 Experiments

Apart from the running example, we present case studies from DSP and control systems to illustrate the utility of the presented synthesis approach. Our technique was implemented in Matlab, and Nelder-Mead implementation available in Matlab as `fminsearch` function was used for numerical optimization. We use the Constantinides et al [6] cost model.

5.1 Running Example

We illustrate the synthesis approach (more details in [14]) presented in Section 4 using the running example. Our algorithm used 34 examples and needed 4 iterations. To evaluate our approach, we exhaustively simulated the generated fixed-point program on the given domain ($0.1 \leq \text{radius} < 2$) at intervals of 0.0001. The is presented in Figure 5.1.




Fig. 7: Our Approach on Running Example.

As a point of comparison, we also show the result of synthesizing a fixed-point program using the `optInduce` routine with 100 inputs (3 times as many as our approach) selected uniformly at random (Figure 5.1). The horizontal line in the plots denotes the maximum error threshold of 0.01 on the relative difference error function. The cost of the fixed-point program synthesized with random sampling is 89.65, and the fixed-point




Fig. 8: Running Example Using Random Inputs.

types of the variables are $\text{fx}\tau(\text{mypi}) = \langle 0, 2, 3 \rangle$, $\text{fx}\tau(\text{radius}) = \langle 0, 1, 8 \rangle$, $\text{fx}\tau(\tau) = \langle 0, 2, 10 \rangle$ and $\text{fx}\tau(\text{area}) = \langle 0, 4, 8 \rangle$. Notice, however, that it is incorrect for a large number of inputs. In contrast, the cost of the implementation produced using our technique is 104.65, and the fixed-point types of the variables are $\text{fx}\tau(\text{mypi}) = \langle 0, 2, 3 \rangle$, $\text{fx}\tau(\text{radius}) = \langle 0, 1, 9 \rangle$, $\text{fx}\tau(\tau) = \langle 0, 2, 11 \rangle$ and $\text{fx}\tau(\text{area}) = \langle 0, 4, 10 \rangle$.

5.2 Infinite Impulse Response (IIR) Filter

The first case study is an IIR filter which is used in digital signal processing applications. It is a first-order direct form-II IIR filter with the schematic shown in Figure 9. The constants are $a_1 = -0.5$, $b_0 = 0.9$ and $b_1 = 0.9$. The fixed-point variables are identified in the schematic. We use our synthesis technique to discover the appropriate fixed-point types of these variables. The input domain used in synthesis is $-2 < \text{input} < 2$. The correctness condition for accuracy is to ensure that the relative error between the floating-point and fixed-point program is less than 0.1.

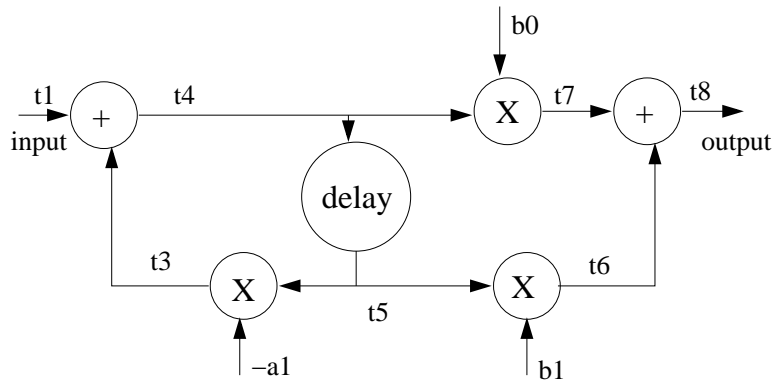


Fig. 9: IIR Filter Schematic




Fig. 10: IIR Filter Using Floating-point and Fixed-point. In the top plot, the floating-point and fixed-point outputs are virtually superimposed on each other.

In order to test the correctness of our implementation, we feed a common input signal to both the IIR filter implementations: floating-point version and the fixed-point version obtained by our synthesis technique. The input signal is a linear chirp from 0 to $\frac{Fs}{2}$ Hz in 1 second.

$$input = (1 - 2^{-15}) \times \sin\left(\pi \times \frac{Fs}{2} \times t^2\right)$$

where $Fs = 256$ and $t = 0$ to $1 - \frac{1}{Fs}$ and is sampled at intervals of $\frac{1}{Fs}$. Figure 10 shows the input, outputs of both implementations and the relative error between the two outputs. We observe that the implementation satisfies the correctness condition and the relative error remains below 0.1 throughout the simulation.

5.3 Finite Impulse Response (FIR) Filter

The second case study is a low pass FIR filter of order 4 with tap coefficients 0.0346, 0.2405, 0.4499, 0.2405 and 0.0346. The input domain, correctness condition and input signal to test the floating-point implementation and synthesized fixed-point program are same as the previous case study. Figure 11 shows the input, outputs of both implementations and the relative error between the two outputs. We observe that the implementation satisfies the correctness condition and the relative error remains below 0.1 throughout the simulation.

5.4 Field Controlled DC Motor

The next case study is a field controlled DC Motor. It is a classic non-linear control example from Khalil [15]. The system dynamics is described by the following ordinary




Fig. 11: FIR Filter Using Floating-point and Fixed-point. The floating-point/fixed-point outputs are virtually superimposed on each other.

differential equations.

$$\begin{aligned}\dot{v}_f &= R_f i_f + L_f \dot{i}_f \\ \dot{v}_a &= c_1 i_f \omega + L_a \dot{i}_a + R_a i_a \\ \dot{J} \omega &= c_2 i_f i_a - c_3 \omega\end{aligned}$$

The first equation is for the field circuit with v_f, i_f, R_f, L_f being its voltage, current, resistance, and inductance. The variables v_a, i_a, R_a, L_a are the corresponding voltage, current, resistance, and inductance of the armature circuit described by the second equation. The third equation is a torque equation for the shaft, with J as the rotor inertia and c_3 as a damping coefficient. The term $c_1 i_f \omega$ is the back electromotive force induced in the armature circuit, and $c_2 i_f i_a$ is the torque produced by the interaction of the armature current with the field circuit flux. In the field controlled DC motor, field voltage v_f is the control input while v_a is held constant. The purpose of the control is to drive the system to the desired set point for the angular velocity ω .

We can now rewrite the system dynamics in the following normal form where $a = \frac{R_f}{L_f}, u = \frac{v_f}{L_f}, b = \frac{R_a}{L_a}, \rho = \frac{c_1}{L_a}, c = \frac{c_2}{L_a}, \theta = \frac{c_2}{J}, d = \frac{c_3}{J}$.

$$\begin{aligned}\dot{i}_f &= -a i_f + u \\ \dot{i}_a &= -b i_a + \rho - c i_f \omega \\ \dot{\omega} &= \theta i_f i_a - d \omega\end{aligned}$$

We assume no damping, that is, $c_3 = 0$ and set all the other constants a, b, c, θ, ρ to 1. The state feedback law to control the system is given by

$$u = \frac{\theta(a+b)i_f i_a + \theta \rho i_f - c \theta i_f^2 \omega}{\theta i_a + \epsilon}$$

where $\epsilon = 0.01$ is added to denominator to avoid division by 0. i_a approaches 0 at equilibrium. The corresponding floating-point code is shown below.

Input: $i_f, i_a, \omega, \theta, \rho, c, \epsilon$

Output: u

```

t1 =  $\theta \times i_a$ ; t2 =  $\epsilon + t1$ ; t3 =  $1/t2$ ; t31 = a + b;
t32 =  $i_f \times i_a$ ; t33 = t31  $\times$  t32; t4 =  $\theta \times t33$ ; t41 =  $\rho \times i_f$ ;
t5 =  $\theta \times t41$ ; t6 =  $i_f \times i_f$ ; t61 = t6  $\times$   $\omega$ ; t62 = t61  $\times$   $\theta$ ;
t7 = c  $\times$  t62; t8 = t4 + t5; t9 = t8 - t7; u = t3  $\times$  t9;
return u

```

The system is initialized with field current $i_f = 1$, armature current $i_a = 1$ and angular velocity $\omega = 1$.




Fig. 12: DC Motor Using Floating-point and Fixed-point Controller with zoomed-in view for 2 to 3 seconds.

The computed control law can be mathematically shown to be correct by designers who are more comfortable in reasoning with real arithmetic but not with finite precision arithmetic. Its implementation using floating-point computation also closely mimics the arithmetic in reals but the control algorithms are often implemented using fixed-point computation on embedded platforms. We use our synthesis technique to automatically derive a low cost fixed-point implementation of the control law computing u . The input domain is $0 \leq i_a, i_f, \omega \leq 1.5$. The correctness condition for accuracy is that the absolute difference between the u computed by fixed-point program and the floating-point program is less than 0.1.

Figure 12 shows the simulation of the system using the fixed-point implementation of the controller and the floating-point implementation. This end-to-end simulation shows that fixed-point program generated by our technique can be used to control the system as effectively as the floating-point program. This illustrates the practical utility of our technique. Figure 13 plots the difference between the control input computed by the fixed-point program and the floating-point program. It shows that the fixed-point types synthesized using our approach satisfy the correctness condition, and the difference between the control input computed by the fixed-point and floating-point program is within the specified maximum error threshold of 0.1. The number of inputs needed in our approach was 127. In contrast, the fixed-point types found using 635(5X our ap-

proach) randomly selected inputs violate the correctness condition for a large number of inputs.




Fig. 13: Error in Control Input Using Fixed-point and Floating-point Program

5.5 Two-Wheeled Welding Mobile Robot

The next case study is a nonlinear controller for a two-wheeled welding mobile robot (WMR) [3]. The robot consists of two wheels and a robotic arm. The wheels can roll and there is no slipping. (x, y) represents the Cartesian coordinate of the WMR's center point and ϕ is the heading angle of the WMR. v and ω are the straight and angular velocities of the WMR at its center point. The welding point coordinates (x_w, y_w) and the heading angle ϕ_w can be calculated from the WMR's center point:

$$\begin{aligned} x_w &= x - l \sin \phi \\ y_w &= y + l \cos \phi \\ \phi_w &= \phi \end{aligned}$$

So, the equation of motion for the welding point is as follows:

$$\begin{aligned} \dot{x}_w &= v \cos \phi - l \omega \cos \phi - \dot{l} \sin \phi \\ \dot{y}_w &= v \sin \phi - l \omega \sin \phi + \dot{l} \cos \phi \\ \dot{\phi}_w &= \omega \end{aligned}$$

The objective of the WMR controller is to ensure that the robot tracks a reference point R. The reference point R moving with a constant velocity of v_r on the reference path has coordinates (x_r, y_r) and the heading angle ϕ_r . The tracking error is the difference between the location of the robot and the reference point.

$$\begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} = \begin{bmatrix} \cos \phi & \sin \phi & 0 \\ -\sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_r - x_w \\ y_r - y_w \\ \phi_r - \phi_w \end{bmatrix}$$

The two control parameters in the model are v and ω . In order to ensure that the error quickly converges to 0, a nonlinear controller based on Lyapunov stability is as follows:

$$v = l(\omega_r + k_2 e_2 v_r + k_3 \sin e_3) + v_r \cos e_3 + k_1 e_1$$

$$\omega = \omega_r + k_2 e_2 v_r + k_3 \sin e_3$$

where k_1, k_2 and k_3 are positive constants. Table 3 provides the numerical values of constants and initial values of the state variables from Bui et al [3]. All lengths are in meters, angle in radians and time in seconds.

Table 3: Numerical and Constant Values

Parameters	Values	Parameters	Values
k_1	4.2	l	0.15
k_2	5000	\dot{i}	0
k_3	1	v_r	$7.5e - 3$
x_r	0.280	x_w	0.270
y_r	0.400	y_w	0.390
ϕ	0	ϕ_w	15




Fig. 14: Reference Welding Line

We use our synthesis technique to automatically synthesize fixed-point program computing both control inputs: v and ω . The error function used for v is the relative difference

$$\frac{v_{floating-point} - v_{fixed-point}}{v_{floating-point}}$$

and the error function used for ω is the moderated relative difference

$$\frac{\omega_{floating-point} - \omega_{fixed-point}}{\omega_{floating-point} + \delta}$$

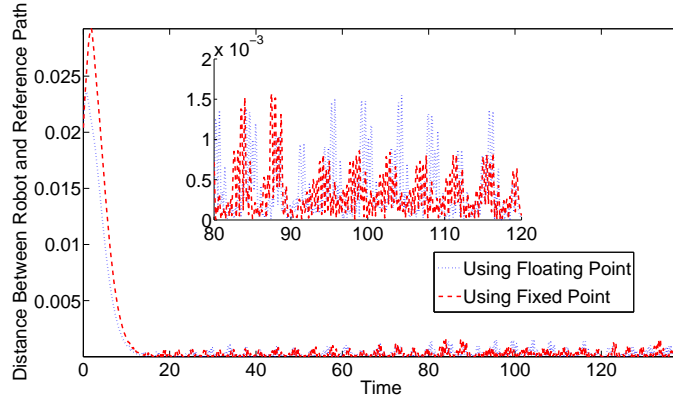


Fig. 15: Distance of WMR from Reference Line with zoomed-in view for 80 to 120 seconds.




Fig. 16: Error in computing v

where $\delta = 0.001$. The moderated relative difference is useful here since $\omega_{floating-point}$ can be 0. We require that the difference values for both controllers are less than 0.1. Figure 14 shows the reference line for welding and Figure 15 shows the distance of the WMR from the reference line as a function of time for both cases: firstly, when the controller is implemented as a floating-point program and secondly, when the controller is implemented as a fixed-point program synthesized using our technique. The robot starts a little away from the reference line but quickly starts tracking the line in both cases. Figure 16 and Figure 17 show the error between the floating-point controller and fixed-point controller for both control inputs: v and ω , respectively.

Table 4 summarizes the performance of our technique in the four case-studies.




Fig. 17: Error in computing ω

Table 4: Performance

Case-study	Runtime (seconds)	# of Iterations
IIR Filter	268	5
FIR Filter	379	4
DC Motor u	4436	8
WMR v	2218	7
ω	1720	4

6 Related Work

Previous techniques for optimizing fixed-point types are based on statistical sampling of the input space. These methods sample a large number of inputs and heuristically solve an optimization problem that minimizes implementation cost while ensuring that some correctness specification is met over the sampled inputs. The techniques differ in the heuristic search method employed, in the measure of cost, or in how accuracy of fixed-point implementation is determined. Sung and Kum [24] use a heuristic search technique which starts with the minimum wordlength implementation as the initial guess. The wordlengths are increased one by one till the error falls below an acceptable threshold. Han et al. [10, 11] use a gradient-based sequential search method which starts with the minimum wordlength implementation as the initial guess. The gradient (ratio of increase in accuracy and increase in wordlengths) is computed for a set of wordlength changes at each step and the search moves in the direction with maximum gradient. Shi et al. [22] propose a floating-point to fixed-point conversion methodology for digital VLSI signal processing systems. Their approach is based on a perturbation theory which shows that the change to the first order is a linear combination of all the first- and second-order statistics of the quantization noise sources. Their technique works with general specification criteria, as long as these can be represented as large ensemble averages of functions of the signal outputs. For example, they use mean-squared error (MSE) as the specification function. The cost of the implementation is a quadratic function. Monte Carlo simulation of a large number of input examples is used to formulate a quadratic optimization problem based on perturbation theory. In contrast, our

specification requires that the accuracy condition holds for all inputs and not just on an average. Further, the cost function can be any arbitrary function for our technique and need not be quadratic. Perhaps most importantly, our technique does not rely on apriori random sampling of a large number of input values, instead using optimization to discover a small set of *interesting* examples which suffice to discover optimal fixed-point implementation. Purely analytical methods [23, 16] based on dataflow analysis have also been proposed for synthesizing fixed-point programs based on forward and backward propagation in the program's dataflow graph. The advantages of these techniques are that they do not rely on picking the right inputs for simulation, can handle arbitrary programs (with approximation), and can provide correctness guarantees. However, they tend to produce very conservative wordlength results. Inductive synthesis based on satisfiability solving has been previously used for synthesizing programs from functional specifications. These approaches [13, 9] rely on constraint solving in much the same way as we rely on optimization routines. However, these approaches only seek to find a correct program, without any notion of cost and optimization. A central intuition in our work is to exploit the nearly continuous nature of error function [20]. These results were first reported in Chapter 4 of [12].

7 Conclusion

In this paper, we presented a novel approach to automated synthesis of fixed-point program from floating-point program by discovering the fixed-point types of the variables. The program is synthesized to satisfy the provided correctness condition for accuracy and to have optimal cost with respect to the provided cost model. We illustrated our approach on a set of case studies from digital signal processing and control systems.

Acknowledgement

This research was done when first author was at UC Berkeley. The research was supported by NSF grants CNS-0644436 and CNS-0627734, the FCRP/MARCO Multi-Scale Systems Center (MuSyC), Microsoft Research, Intel, and the Berkeley Fellowship for Graduate Studies from UC Berkeley.

References

1. ANSI/IEEE Std. 754-1985: IEEE standard for binary floating-point arithmetic.
2. Matlab: Fixed-point toolbox.
3. T.H. Bui, T.T. Nguyen, T.L. Chung, and S.B. Kim. A simple nonlinear control of a two-wheeled welding mobile robot. *Int. Journal of Control, Automation and Systems*, 1, 2003.
4. Swarat Chaudhuri, Sumit Gulwani, and Roberto Lubliner. Continuity analysis of programs. In *POPL '10*, pages 57–70, 2010.
5. Jonathan A. Clarke, Altaf Abdul Gaffar, George A. Constantinides, and Peter Y. K. Cheung. Fast word-level power models for synthesis of FPGA-based arithmetic. In *ISCAS*, 2006.
6. George A. Constantinides, Peter Y. K. Cheung, and Wayne Luk. Wordlength optimization for linear digital signal processing. In *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol 22, No. 10, pages 1432–1443, 2003.
7. R. Fletcher. *Practical Methods of Optimization*. J. Wiley, 1986.

8. David Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Computing Surveys*, 23:5–48, 1991.
9. Sumit Gulwani, Susmit Jha, Ashish Tiwari, and Ramarathnam Venkatesan. Synthesis of loop-free programs. In *PLDI*, pages 62–73, 2011.
10. Kyungtae Han, Iksu Eo, Kyungsu Kim, and Hanjin Cho. Numerical word-length optimization for cdma demodulator. In *ISCAS 2001*, volume 4, pages 290–293 vol. 4, may 2001.
11. Kyungtae Han and B.L. Evans. Wordlength optimization with complexity-and-distortion measure and its application to broadband wireless demodulator design. In *IEEE International Conference on Acoustics, Speech, and Signal Processing, 2004*, volume 5, pages V – 37–40 vol.5, may 2004.
12. Susmit Jha. *Towards Automated System Synthesis Using SCIDUCTION*. PhD thesis, EECS Department, University of California, Berkeley, Nov 2011.
13. Susmit Jha, Sumit Gulwani, Sanjit Seshia, and Ashish Tiwari. Oracle-guided component-based program synthesis. In *ICSE*, pages 215–224, 2010.
14. Susmit Jha and Sanjit A. Seshia. SWATI: Synthesizing wordlengths automatically using testing and induction. *UC Berkeley Technical Report*, pages TR–XX, 2013.
15. H.K. Khalil. *Nonlinear systems*. Macmillan Pub. Co., 1992.
16. Seehyun Kim, Ki-II Kum, and Wonyong Sung. Fixed-point optimization utility for c and c++ based digital signal processing programs. *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, 45(11):1455–1464, nov 1998.
17. Mike Tien-Chien Lee, V. Tiwari, S. Malik, and M. Fujita. Power analysis and minimization techniques for embedded DSP software. *VLSI*, 5(1):123–135, march 1997.
18. E. Macii, M. Pedram, and F. Somenzi. High-level power modeling, estimation, and optimization. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 17(11):1061–1079, nov 1998.
19. Rupak Majumdar and Indranil Saha. Symbolic robustness analysis. In *RTSS '09*, pages 355–363, 2009.
20. Rupak Majumdar, Indranil Saha, and Majid Zamani. Synthesis of minimal-error control software. In *EMSOFT*, pages 123–132, 2012.
21. J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, January 1965.
22. Changchun Shi and Robert W. Brodersen. An automated floating-point to fixed-point conversion methodology. In *Proc. IEEE Int. Conf. on Acoust., Speech, and Signal Processing*, pages 529–532, 2003.
23. Mark Stephenson, Jonathan Babb, and Saman Amarasinghe. Bidwidth analysis with application to silicon compilation. *SIGPLAN Not.*, 35:108–120, May 2000.
24. Wonyong Sung and Ki-II Kum. Simulation-based word-length optimization method for fixed-point digital signal processing systems. *Signal Processing, IEEE Transactions on*, 43(12):3087–3090, dec 1995.
25. Randy Yates. Fixed-point arithmetic. In *Technical Reference Digital Signal Labs*, 2009.