

Verifying Quantitative Reliability of Programs That Execute on Unreliable Hardware (Appendix)

Michael Carbin Sasa Misailovic Martin C. Rinard

MIT CSAIL

{mcarbin, misailo, rinard}@csail.mit.edu

A. Additional Semantics Definitions

This section contains additional semantics definitions elided from the presentation in Section 3 in the paper.

A.1 Semantics of Arrays

Figure 11 presents the dynamic semantics of array operations.

Declarations. An array declaration allocates a new array in the heap. The boundaries of an array are given by a sequence of expressions, each of which can evaluate unreliably [E-ARRAY-DECL-R]. Given these boundaries, a new array is allocated using the function *new*, which returns the base address of an array that has been freshly allocated in the memory region m [E-ARRAY-DECL]. The function *new* executes reliably. To guard against ill-defined behavior, the given semantics also reliably checks that the length of each dimension is non-negative.

Loads. Executing an array load entails multiple steps. In the first step, the program reduces the expressions for each index in left-to-right order [E-ARRAY-LOAD-IDX]. Note that reduction of the index expressions may encounter faults, producing incorrect indices. While our dynamic semantics incorporates bounds checks to prevent incorrect indices from yielding ill-defined, out-of-bounds behaviors, we have left the choice of index expression reliability to the discretion of the developer.

Given the array reference and the reduced indices of each dimension, the program then checks to see if each index is within bounds of the allocated dimensions of the array and also calculates the offset of the element to be accessed [E-ARRAY-LOAD-C, E-ARRAY-LOAD-F]. In the final step, the array load proceeds by attempting to fetch the corresponding value from the given memory region. With probability $\psi(rd(m))$, this step executes correctly and returns the value from memory [E-ARRAY-LOAD-C]. With probability $1 - \psi(rd(m))$, the memory read fails, producing an alternative value n_f with probability $P(n_f)$ [E-ARRAY-LOAD-F].

Stores. The semantics of array stores are similar to that for loads except with the reliability of writes to the array’s memory region ($\psi(wr(m))$) substituted for the reliability of

reads. We have elided a full presentation of the semantics of array stores for brevity.

We note that the store operation may fail in two ways: 1) the index computation produces a wrong index (the rule is similar to [E-ARRAY-LOAD-IDX]) or 2) the write operation may fail, with probability $1 - \psi(wr(m))$. If the index computation fails and the computed index is outside of the loop bounds, then the computation skips the write operation. If the index computation fails and the computed index is within the bound, then the write will modify another location within the array. If the write operation fails, the semantics stores an alternative value n_f with probability $P(n_f)$.

A.2 Semantics and Analysis of Functions

A.2.1 Dynamic Semantics

Figure 12 presents the dynamic semantics of function calls and returns. A function call and return sequence executes via the following procedure:

Call Argument Evaluation. A function call first evaluates its arguments in left-to-right order [E-CALLX-ARGS]. Note that evaluation of the arguments may encounter faults. A parameter value v is either value that a numerical expression evaluates to, which is denoted as v^{num} or a memory location of the array, which is denoted as v^{ref} .

Call Body Unfolding. After the arguments to a call have been fully evaluated, control then transfers to the body of function [E-CALLX-UNFOLD]. The rule transfers control by fetching the code for the body of the function via the utility function *code(f)*.

The rule also creates a new state σ' and pushes the old state σ onto the program stack δ . The rule initializes the new state with the appropriate values for its parameters by prepending a sequence of declarations *s_{int.init}* to allocate and initialize the numerical parameters of the function. The function *array_params* returns the set of names and position of the formal array parameters of the function f . The function *int_params* returns the set of names, memory locations, and positions of integer formal parameters of the function f .

Note that this rule executes correctly with probability 1; this implies that both control transfers and manipulations of the program stack are performed fully reliably. While

E-ARRAY-DECL-R

$$\frac{\langle e_i, \sigma \rangle \xrightarrow{\theta, p}_\psi e'_i}{\langle \mathbf{int} \ a [n_1, \dots, e_i, \dots, e_k] \ \mathbf{in} \ m, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle \mathbf{int} \ a [n_1, \dots, e'_i, \dots, e_k] \ \mathbf{in} \ m, \langle \sigma :: \delta, h \rangle \rangle}$$

E-ARRAY-DECL

$$\frac{\forall i. 0 < n_i \quad \langle n_b, h' \rangle = \mathit{new}(h, m, \langle n_1, \dots, n_k \rangle) \quad \sigma' = \sigma[a \mapsto \langle n_b, \langle n_1, \dots, n_k \rangle, m \rangle]}{\langle \mathbf{int} \ a [n_1, \dots, n_k] \ \mathbf{in} \ m, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{C, 1}_\psi \langle \mathbf{skip}, \langle \sigma' :: \delta, h' \rangle \rangle}$$

E-ARRAY-LOAD-IDX

$$\frac{\langle e_i, \sigma \rangle \xrightarrow{\theta, p}_\psi e'_i}{\langle x = a [n_1, \dots, e_i, \dots, e_k], \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle x = a [n_1, \dots, e'_i, \dots, e_k], \langle \sigma :: \delta, h \rangle \rangle}$$

E-ARRAY-LOAD-C

$$\frac{\sigma(a) = \langle n_b, \langle l_1, \dots, l_k \rangle, m \rangle \quad n_o = l_k + \sum_{i=0}^{k-1} n_i \cdot l_i \quad n = h(n_b + n_o)}{\langle x = a [n_1, \dots, n_k], \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{C, \psi(rd(m))}_\psi \langle x = n, \langle \sigma :: \delta, h \rangle \rangle}$$

E-ARRAY-LOAD-F

$$\frac{\sigma(a) = \langle n_b, \langle l_1, \dots, l_k \rangle, m \rangle \quad n_o = l_k + \sum_{i=0}^{k-1} n_i \cdot l_i \quad p = (1 - \psi(rd(m))) * P(n_f \mid rd(m))}{\langle x = a [n_1, \dots, n_k], \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{(F, n_f), p}_\psi \langle x = n_f, \langle \sigma :: \delta, h \rangle \rangle}$$

Figure 11: Dynamic Semantics of Arrays

E-CALLX-EXPR-ARG

$$\frac{\langle e_i, \sigma \rangle \xrightarrow{\theta, p}_\psi e'_i}{\langle x = f(v_1, \dots, e_i, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle x = f(v_1, \dots, e'_i, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle}$$

E-CALLX-INT-ARG

$$\frac{\langle e_i, \sigma \rangle \xrightarrow{\theta, p}_\psi n \quad v_i^{num} = n}{\langle x = f(v_1, \dots, e_i, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle x = f(v_1, \dots, v_i^{num}, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle}$$

E-CALLX-ARR-ARG

$$\frac{v_i^{ref} = \sigma(a_i)}{\langle x = f(v_1, \dots, a_i, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{C, 1}_\psi \langle x = f(v_1, \dots, v_i^{ref}, \dots, e_k), \langle \sigma :: \delta, h \rangle \rangle}$$

E-CALLX-UNFOLD

$$\frac{\forall (a_i, i) \in \mathit{array_params}(f) . \sigma'(a_i) = v_i^{ref} \quad s_{\mathit{int_init}} \equiv \mathbf{int} \ x_{i_1} = v_{i_1}^{num} \ \mathbf{in} \ m_{i_1} ; \dots ; \mathbf{int} \ x_{i_k} = v_{i_k}^{num} \ \mathbf{in} \ m_{i_k} \ \mathbf{where} \ (x_{i_j}, m_{i_j}, i_j) \in \mathit{int_params}(f)}{\langle x = f(v_1, \dots, v_k), \langle \delta, h \rangle \rangle \xrightarrow{C, 1}_\psi \langle x = f(v_1, \dots, v_k) \ s_{\mathit{int_init}} ; \mathit{code}(f), \langle \sigma' :: \delta, h \rangle \rangle}$$

E-CALLX-BODY

$$\frac{\langle s, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle s', \langle \sigma' :: \delta', h' \rangle \rangle}{\langle x = f(v_1, \dots, v_k) \ s, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle x = f(v_1, \dots, v_k) \ s', \langle \sigma' :: \delta', h' \rangle \rangle}$$

E-RETURN-R

$$\frac{\langle e, \sigma \rangle \xrightarrow{\theta, p}_\psi \langle e', \sigma \rangle}{\langle \mathbf{return} \ e, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{\theta, p}_\psi \langle \mathbf{return} \ e', \langle \sigma :: \delta, h \rangle \rangle}$$

E-CALLX-RETURN

$$\frac{}{\langle x = f(v_1, \dots, v_k) \ \mathbf{return} \ n, \langle \sigma :: \delta, h \rangle \rangle \xrightarrow{C, 1}_\psi \langle x = n, \langle \delta, h \rangle \rangle}$$

E-SEQ-RETURN

$$\frac{s_1 \in \{\mathbf{return}, \mathbf{return} \ n\}}{\langle s_1 ; s_2, \langle \delta, h \rangle \rangle \xrightarrow{C, 1}_\psi \langle s_1, \langle \delta, h \rangle \rangle}$$

Figure 12: Dynamic Semantics of Function Calls and Returns

$$\begin{aligned}
RP_\psi(x = f(v_1, \dots, v_n), Q, C) &= Q[\psi(\text{wr}(\Lambda(x)) \cdot \text{Relspec_act}(f, 0, X) / \mathcal{R}(\{x\} \cup X)) \\
&\quad [\text{Relspec_act}(f, i_1, X) / \mathcal{R}(\{a_{i_1}\} \cup X)] \dots \\
&\quad [\text{Relspec_act}(f, i_k, X) / \mathcal{R}(\{a_{i_k}\} \cup X)] \\
\text{Relspec_act}(f, i, X) &= \text{Relspec_form}(f, i) \\
&\quad [\rho_1(\text{act_par}(f, 1)) \cdot \mathcal{R}(\rho_2(\text{act_par}(f, 1)) \cup Y) / \mathcal{R}(\{\text{form_par}(f, 1)\} \cup Y)] \dots \\
&\quad [\rho_1(\text{act_par}(f, j)) \cdot \mathcal{R}(\rho_2(\text{act_par}(f, j)) \cup Y) / \mathcal{R}(\{\text{form_par}(f, j)\} \cup Y)] \dots \\
&\quad [\rho_1(\text{act_par}(f, n-1)) \cdot \mathcal{R}(\rho_2(\text{act_par}(f, n-1)) \cup Y) / \mathcal{R}(\{\text{form_par}(f, n-1)\} \cup Y)] \\
&\quad [\rho_1(\text{act_par}(f, n)) \cdot \mathcal{R}(\rho_2(\text{act_par}(f, n)) \cup X) / \mathcal{R}(\{\text{form_par}(f, n)\})]
\end{aligned}$$

Figure 13: Constraint Generation for Function Calls

reliable control transfers are given by Rely’s machine model, reliable program stack manipulations require that a compiler allocate the program stack in a reliable memory region.

Return Value Evaluation. A return e statement fully evaluates e under the hardware reliability model, yielding a value n [E-RETURN-E].

Call Return Execution. Once the body of the call executes and reaches a return statement, execution proceeds by restoring the old state σ' from the program stack, assigning the return value to the destination variable, and transferring control back to the caller [E-CALL-RETURN-E]. Note that as with [E-CALLX-UNFOLD], this step executes correctly with probability 1 and therefore both the control transfer and stack manipulation are fully reliable.

B. Constraint Generation for Functions

Figure 13 presents the constraint generation rule for function calls. Intuitively, for a function call $x = f(e_1, \dots, e_n)$ the constraint generator performs two tasks. First, it substitutes the declared reliabilities of with the reliability expressions for the actual parameters of the function known at the call site. Second, to update the reliability expressions for the modified array variables, it constructs a constraint that abstracts the reliability of a function call as the reliability of multiple assignment statements – one statement represents the assignment of the final value of the function to the variable x , the remaining statements represent the assignment of each potentially modified array parameter of the function.

To substitute the return value, the rule in Figure 13 computes the reliability expression derived from the reliability specification of the function’s return value. The helper function *Relspec_act* replaces substitutes the reliability of all actual parameters of the function’s call for the declared parameter names within the declared reliability expression for the return value (when $i = 0$) or one of the function’s parameters (when $i \geq 1$). The rule then identifies array variables (in the presentation there is total of $k \leq n$ array variables) that are passed as parameters to the function and updates their reliabilities analogously.

The helper function *Relspec_act* takes as input the function’s specification, the index i of the array parameter or 0

for the return value, and the set of the variables from the context of the calling function. The function 1) obtains the reliability specification associated with the i -th parameter of the function, 2) the name of the i -th formal parameter using the *formal_arg* function, and 3) the reliability of the actual parameter expression using the *act_arg* function. The function *act_arg* returns the reliability of an expression if the actual parameter is a numerical value or the reliability pair $(1.0, \mathcal{R}(a_i))$ if the actual parameter is an array variable.

Then the function *Relspec_act* substitutes reliabilities of all function’s formal parameters with the expressions of the corresponding actual parameters. The resulting expression does not contain the names of the formal parameters. We note that the rule for function calls requires that the names of the formal parameters are distinct from the names of variables used in the actual parameter reliability expressions. This can be enforced by additional renaming of formal parameters.

C. Properties of Reliability Factors

In this section we present proofs of several properties of reliability factors.

C.1 Discrete Distribution

This theorem ensures that the distribution of unreliable distributions is discrete, i.e., it can be represented by a probability mass function φ .

Lemma 1 (Discrete Distribution). *The probability space of unreliable environments (E, φ) is discrete.*

Proof. A probability distribution is discrete if it is defined on a countable sample space. Therefore, we need to check that the set E is countable. In this proof we denote the cardinality of a set X as $|X|$ and as a special case the cardinality of a set of natural numbers, $|\mathbb{N}| = \aleph_0$.

The cardinality of the set of heaps is $|\text{Loc} \times \text{Int}_M|$. Since both $|\text{Loc}|$ and $|\text{Int}_M|$ are finite, the cardinality of their product is also finite. The cardinality of a single stack frame is $|\Sigma| = |\text{Var} \times \text{Ref}|$. The cardinality of Var is finite since a function can have only a finite number of variables. Cardinality $|\text{Ref}| = |\text{Loc} \times \text{Loc}^k \times M|$ is finite, since Loc and M are finite, and the number of the array dimensions is a finite

number (which follows from the definition of the program's syntax). The cardinality of a stack is then equal to \aleph_0 since as a finite or a countable length list of finite-length frames, it can be mapped onto \mathbb{N} . The cardinality of E is equal to the product of cardinalities of $|H|$ and $|\Sigma|$, and is equal to \aleph_0 . Therefore, E is a countable set. \square

C.2 Reliability Factor Ordering

The ordering proposition states that the function $\mathcal{R}(X)$ is monotonically decreasing – for a larger set of variables X , the probability that the variables in that set have the same value decreases. This proposition also enables comparison of symbolic reliability predicates.

Proposition 1 (Ordering; restated). *For two sets of variables X and Y , if $X \subseteq Y$ then $\mathcal{R}(Y) \leq \mathcal{R}(X)$.*

Proof. First, we consider the case when all variables in X and Y are scalars. Let $U_Y = \mathcal{E}(Y, \varepsilon)$ be the set of alternative environments in which all variables in the set Y have the same value as in the original environment and let $U_X = \mathcal{E}(X, \varepsilon)$ be the set of alternative environments in which all variables in the set X have the same value as in the original environment. Then, if $X \subseteq Y$, the set $U_Y \subseteq U_X$, since the variables in $Y \setminus X$ provide additional restrictions on the states that are contained in U_Y . Finally, the theorem statement follows from the inequality $\sum_{v \in U_X} \varphi(v) \geq \sum_{v \in U_Y} \varphi(v)$.

If a is an array variable, then the function *equiv* adds a constraint for each element of a . Then, we can apply the same argument for each such obtained sets U_X and U_Y . \square

As a side note, joint reliability of multiple variables can be obtained from the marginal reliabilities of these variables:

Proposition 1 (Union Bound). *Let $X = \{x_1, \dots, x_n\}$ be a set of n program variables. Then, $\mathcal{R}(X) \geq \sum_{i=1}^n \mathcal{R}(\{x_i\}) - (n - 1)$.*

Proof. The probability that the in the set of variables X , at least one variable has an incorrect value is $1 - \mathcal{R}(X)$. The probability that a variable x_i has an incorrect value is $1 - \mathcal{R}(\{x_i\})$.

The probability that any of variables in X has an incorrect value is smaller than the probability that every individual program variable has an incorrect value (from the classical probabilistic union bound). Then, the statement follows from $1 - \mathcal{R}(X) \leq \sum_{i=1}^n 1 - \mathcal{R}(\{x_i\})$. \square

C.3 Redundant Constraints

As a consequence of ordering we can define the simplification procedure that removing redundant predicates from the constraints that the analysis produces.

Proposition 2 (Predicate Subsumption). *A reliability predicate $r_1 \cdot \mathcal{R}(X_1) \leq r_2 \cdot \mathcal{R}(X_2)$ subsumes (i.e., soundly replaces) a predicate $r'_1 \cdot \mathcal{R}(X'_1) \leq r'_2 \cdot \mathcal{R}(X'_2)$ if $r'_1 \cdot \mathcal{R}(X'_1) \leq r_1 \cdot \mathcal{R}(X_1)$ and $r_2 \cdot \mathcal{R}(X_2) \leq r'_2 \cdot \mathcal{R}(X'_2)$*

Proof. If $P = r_1 \cdot \mathcal{R}(X_1) \leq r_2 \cdot \mathcal{R}(X_2)$ and $P' = r'_1 \cdot \mathcal{R}(X'_1) \leq r'_2 \cdot \mathcal{R}(X'_2)$, we show that $P \wedge P_{rest} \implies P \wedge P' \wedge P_{rest}$, where P_{rest} are the remaining predicates in the constraint.

Let us first consider the case when P_{rest} is empty. Then, we need to show that $P \implies P \wedge P'$. This predicate follows from the inequality assumptions from the theorem statement. Specifically, if $r_1 \cdot \mathcal{R}(X_1) \leq r_2 \cdot \mathcal{R}(X_2)$ is true, by the ordering lemma and the widening of the interval, $r'_1 \cdot \mathcal{R}(X'_1) \leq r'_2 \cdot \mathcal{R}(X'_2)$ is also true.

If P_{rest} is not empty, one can replace the logical value true or false for P_{rest} and then use the same argument as in the previous derivation. \square

```

1 #define tolerance 0.000001
2 #define maxsteps 40
3
4 float<0.9999*R(x)> F(float x in unrel);
5
6 float<0.9999*R(x)> dF(float x in unrel);
7
8 float <0.995*R(xa, xb)> secant
9 (float xa in unrel, float xb in unrel) {
10     float a in unrel;
11     float b in unrel;
12     float c in unrel;
13     float fa in unrel;
14     float fb in unrel;
15     float fc in unrel;
16     bool converged in unrel;
17
18     a = xa;
19     b = xb;
20
21     fa = F(a);
22     fb = F(b);
23
24     while ((a -. b >=. tolerance) ||
25           (a -. b <=. 0.0-.tolerance))
26         : maxsteps {
27         c = (a +. b) /.2;
28         fc = F(c);
29
30         if ((fa >= 0) &&. (fc >= 0)) {
31             a = c;
32             fa = fc;
33         } else {
34             b = c;
35             fb = fc;
36         }
37     }
38     if (!(a -. b <=. tolerance) &&.
39         (a -. b >=. 0.0-.tolerance))) {
40         x = INFTY;
41     }
42     return x;
43 }

```

Figure 14: Secant Method Implementation

```

1 //Array with sine polynomial coefficients
2 #define nsin 20
3 const csin (1) = { /*...*/ };
4
5 //Array with cosine polynomial coefficients
6 #define ncos 19
7 const ccos (1) = { /*...*/ };
8
9 float<0.99999*R(x)> usin(float x in unrel) {
10     float res in unrel;
11     float t in unrel;
12
13     int i; i = 1;
14
15     res = csin[0];
16     while (true) : nsin {
17         t = csin[i];
18         res = res *. x +. t;
19         i = i + 1;
20     }
21     return res;
22 }
23
24 float<0.99999*R(x)> ucos(float x in unrel) {
25     float res in unrel;
26     float t in unrel;
27
28     int i;
29     i = 1;
30
31     res = ccos[0];
32     while (true) : ncos {
33         t = ccos[i];
34         res = res *. x +. t;
35         i = i + 1;
36     }
37     return res;
38 }
39
40 void main
41 (float r in unrel, float theta in unrel
42  float<0.99995*R(r, theta, xy)> xy) {
43
44     float x in unrel;
45     float y in unrel;
46     float t in unrel;
47
48     t = ucos(theta);
49     xy(1) = r *. t;
50     t = usin(theta);
51     xy(2) = r *. t;
52
53 }
54 }

```

Figure 15: Coordinate Conversion Implementation

```

1  const num matx = 64;
2  const num maty = 64;
3
4  void matvec (float mat(2) in urel,
5              float v(1) in urel,
6              num<.997*R(mat, vec, outvec)>
7              u(1) in urel)
8  {
9      num t1 in urel;
10     num t2 in urel;
11     num t3 in urel;
12     int i, j, k;
13
14     i = 0;
15     repeat matx {
16         j = 0;
17         t3 = 0;
18         repeat maty {
19             t1 = mat[j, i];
20             t2 = v[j];
21             t3 = t3 +. t1 *. t2;
22         }
23
24         u[i] = t3;
25         i = i + 1;
26     }
27 }

```

Figure 16: Matrix-Vector multiplication Implementation

D. Benchmarks

D.1 Newton's Method

This is the computation we presented in Section 6.2.

D.2 Secant Method

This computation also searches for a root of function f . It takes as its input two points x_a and x_b for which the function has the opposite sign and returns the value x_0 within the interval $[x_a, x_b]$ for which the function f evaluates to zero. The result of the computation can be checked (as in Newton's method) by ensuring that $f(x_0)$ is close to zero.

This is the fixed point computation that at each step computes the middle point x_c and its function $f(x_c)$. Based on the sign of $f(x_c)$ the algorithm divides the search interval (using a conditional statement). The whole loop can execute unreliably. If the function f has the reliability specification $\text{float}<0.9999*R(x)> F(\text{float } x)$, then the analysis verifies that the reliability of the computation is at least $.995*R(x_a, x_b)$.

Figure 14 presents the implementation of this computation. Note that the reliability of this computation is higher than the reliability of Newton's method. This is because the Secant method makes a single call to the function f , whereas Newton's method makes calls to both f and f' in every iteration. Given multiple options (such as Newton's method and Secant), verified reliability specifications may help developers select the option that best satisfies their combined reliability and efficiency goals.

D.3 Coordinate Conversion

This computation converts polar coordinates of a point to Cartesian coordinates. Given a coordinate (R, θ) , where R is a radius and θ is an angle, it computes the coordinates $x = R \cos(\theta)$ and $y = R \sin(\theta)$. This is also a checkable computation: the square of the diameter is equal to the sum of the squares of Cartesian coordinates.

Figure 15 presents the implementation of this computation. The function takes the inputs `r` and `theta` and stores the result in the output array `xy`. The body of the computation can execute unreliably. The analysis first verifies the implementation of the trigonometric functions (which are evaluations of appropriate Chebyshev interpolating polynomials). If the reliability of the input x is $R(x)$, then the analysis verifies that the functions $\sin(x)$ and $\cos(x)$ have the reliability at least $0.99999*R(x)$. The analysis uses this result to verify that if the reliability of the parameters `r` and `theta` is $R(r, \text{theta})$, then the reliability of the coordinate conversion computation is at least $0.99995*R(r, \text{theta}, xy)$.

D.4 Motion Estimation

This is the computation we presented in Section 2.

D.5 Matrix-vector Multiplication

This computation calculates the product of a matrix M and a vector v . The dimensions of the matrix are $w \times h$, the length of v is h , and the length of the resulting vector u is w . The computation takes `M`, `v`, and `u` as inputs, and computes the values of the elements of the vector `u`. This computation is an important kernel in numerical computations. For example, a linear operator (such as Fourier or Discrete Cosine transforms of arbitrary sizes) can be implemented as a matrix-vector multiplication.

Figure 16 presents the implementation of this computation. All operations on the input data can execute unreliably. The output of the function is the vector `u`. Assuming the maximum size of the square matrix to be 64×64 (as in some signal processing applications), the specified output reliability for the vector `u` is $0.997*R(M, v, u)$. The analysis result states that the reliability of *any* element of `u` is greater than this specified output reliability. Note that the specification needs to account for the possible unreliability of assignments to the vector `u` before entering the function, because it does not track the array indices to determine that every element of `u` is modified inside the function.

D.6 Hadamard Transform

This computation takes as input two blocks of 4×4 pixels and computes the sum of differences between the pixels in the frequency domain. This computation is used in digital signal processing applications (including motion estimation algorithms).

Figure 17 presents the implementation of this computation. The computation calculates the intermediate distances

```

1  int <0.9999995 * R(val)> abs
2    (int val in unrel)
3  {
4    int t in unrel = val;
5    if (t <= . 0) {
6      t = 0 -. t;
7    }
8    return t;
9  }
10
11 int<0.99995 * R(bA, bB, satdstart)>
    hadamarddiff
12   (int<R(bA)> bA(2) in unrel,
13    int<R(bB)> bB(2) in unrel,
14    int satdstart in unrel)
15 {
16   int isatd in unrel;
17   int tmp00 in unrel; int tmp01 in unrel;
18   int tmp02 in unrel; int tmp03 in unrel;
19   int tmp10 in unrel; int tmp11 in unrel;
20   int tmp12 in unrel; int tmp13 in unrel;
21   int tmp20 in unrel; int tmp21 in unrel;
22   int tmp22 in unrel; int tmp23 in unrel;
23   int tmp30 in unrel; int tmp31 in unrel;
24   int tmp32 in unrel; int tmp33 in unrel;
25   int t1 in unrel; int t2 in unrel;
26   int t3 in unrel; int t4 in unrel;
27   int tt in unrel;
28
29   t1 = bA[0,0]; tt = bB[0,0]; t1 = t1 -. tt;
30   t2 = bA[0,1]; tt = bB[0,1]; t2 = t2 -. tt;
31   t3 = bA[0,2]; tt = bB[0,2]; t3 = t3 -. tt;
32   t4 = bA[0,3]; tt = bB[0,3]; t4 = t4 -. tt;
33   tmp00 = t1 +. t2 +. t3 +. t4;
34   tmp01 = t1 +. t2 -. t3 -. t4;
35   tmp02 = t1 -. t2 -. t3 +. t4;
36   tmp03 = t1 -. t2 +. t3 -. t4;
37
38   t1 = bA[1,0]; tt = bB[1,0]; t1 = t1 -. tt;
39   t2 = bA[1,1]; tt = bB[1,1]; t2 = t2 -. tt;
40   t3 = bA[1,2]; tt = bB[1,2]; t3 = t3 -. tt;
41   t4 = bA[1,3]; tt = bB[1,3]; t4 = t4 -. tt;
42   tmp10 = t1 +. t2 +. t3 +. t4;
43   tmp11 = t1 +. t2 -. t3 -. t4;
44   tmp12 = t1 -. t2 -. t3 +. t4 ;
45   tmp13 = t1 -. t2 +. t3 -. t4 ;
46
47   t1 = bA[2,0]; tt = bB[2,0]; t1 = t1 -. tt;
48   t2 = bA[2,1]; tt = bB[2,1]; t2 = t2 -. tt;
49   t3 = bA[2,2]; tt = bB[2,2]; t3 = t3 -. tt;
50   t4 = bA[2,3]; tt = bB[2,3]; t4 = t4 -. tt;
51   tmp20 = t1 +. t2 +. t3 +. t4;
52   tmp21 = t1 +. t2 -. t3 -. t4;
53   tmp22 = t1 -. t2 -. t3 +. t4;
54   tmp23 = t1 -. t2 +. t3 -. t4;
55
56   t1 = bA[3,0]; tt = bB[3,0]; t1 = t1 -. tt;
57   t2 = bA[3,1]; tt = bB[3,1]; t2 = t2 -. tt;
58   t3 = bA[3,2]; tt = bB[3,2]; t3 = t3 -. tt;
59   t4 = bA[3,3]; tt = bB[3,3]; t4 = t4 -. tt;
60   tmp30 = t1 +. t2 +. t3 +. t4;
61   tmp31 = t1 +. t2 -. t3 -. t4;
62   tmp32 = t1 -. t2 -. t3 +. t4;
63   tmp33 = t1 -. t2 +. t3 -. t4;
64
65   isatd = satdstart;
66
67   t1 = abs(tmp00 +. tmp10 +. tmp20 +. tmp30);
68   t2 = abs(tmp00 +. tmp10 -. tmp20 -. tmp30);
69   t3 = abs(tmp00 -. tmp10 -. tmp20 +. tmp30);
70   t4 = abs(tmp00 -. tmp10 +. tmp20 -. tmp30);
71   isatd = isatd +. t1 +. t2 +. t3 +. t4;
72
73   t1 = abs(tmp01 +. tmp11 +. tmp21 +. tmp31);
74   t2 = abs(tmp01 +. tmp11 -. tmp21 -. tmp31);
75   t3 = abs(tmp01 -. tmp11 -. tmp21 +. tmp31);
76   t4 = abs(tmp01 -. tmp11 +. tmp21 -. tmp31);
77   isatd = isatd +. t1 +. t2 +. t3 +. t4;
78
79   t1 = abs(tmp02 +. tmp12 +. tmp22 +. tmp32);
80   t2 = abs(tmp02 +. tmp12 -. tmp22 -. tmp32);
81   t3 = abs(tmp02 -. tmp12 -. tmp22 +. tmp32);
82   t4 = abs(tmp02 -. tmp12 +. tmp22 -. tmp32);
83   isatd = isatd +. t1 +. t2 +. t3 +. t4;
84
85   t1 = abs(tmp03 +. tmp13 +. tmp23 +. tmp33);
86   t2 = abs(tmp03 +. tmp13 -. tmp23 -. tmp33);
87   t3 = abs(tmp03 -. tmp13 -. tmp23 +. tmp33);
88   t4 = abs(tmp03 -. tmp13 +. tmp23 -. tmp33);
89   isatd = isatd +. t1 +. t2 +. t3 +. t4;
90   return isatd;
91 }

```

Figure 17: Hadamard Transform Implementation

between the elements stored in the arrays and computes the sum of the absolute differences of combinations of these elements. The video can be stored in unreliable memory and the entire computation can execute unreliably. The analysis ver-

ifies that if the reliability of the two input blocks bA and bB is $R(bA, bB)$, then the reliability of the computation is greater than $0.99995 * R(bA, bB)$.