

Diamont: Dynamic Monitoring of Uncertainty for Distributed Asynchronous Programs

Vimuth Fernando, Keyur Joshi, Jacob Laurel, and Sasa Misailovic

University of Illinois Urbana-Champaign
{wvf2, kpjoshi2, jlaurel2, misailo}@illinois.edu

Abstract. Many application domains including graph analytics, the Internet-of-Things, precision agriculture, and media processing operate on noisy data and/or produce approximate results. These applications can distribute computation across multiple (often resource-constrained) processing units. Analyzing the reliability and accuracy of such applications is challenging, since most existing techniques operate on specific fixed error models, check for individual properties, or can only be applied to sequential programs. We present Diamont, a system for dynamic monitoring of uncertainty properties in distributed programs. Diamont programs consist of distributed processes that communicate via asynchronous message passing. Diamont includes datatypes that dynamically monitor uncertainty in data and provides support for checking predicates over the monitored uncertainty at runtime. We also present a general methodology for verifying the soundness of the runtime system and optimizations using canonical sequentialization. We implemented Diamont for a subset of the Go language and evaluated eight programs from precision agriculture, graph analytics, and media processing. We show that Diamont can prove important end-to-end properties on the program outputs for significantly larger inputs compared to prior work, with modest execution time overhead: 3% on average and 16.3% at maximum.

1 Introduction

Many emerging distributed applications operate on inherently noisy data or produce approximate results [41]. Emerging edge applications, including autonomous robotics and precision agriculture, routinely need to deal with noise from their sensors. Machine learning applications regularly encounter datasets that contain a high degree of noise, or other irregularity. Furthermore, the rise of highly-parallel and often heterogeneous systems have brought forth new challenges in overcoming bottlenecks in computation and communication between processing units. Many prominent systems adopted approximation in communication, e.g., MapReduce’s task dropping [16], TensorFlow’s precision reduction [43], or Hogwild’s synchronization-eschewing stochastic gradient descent [31]. Also, researchers explored various non-conventional architectures and networks-on-chip [7, 17, 30, 42].

To cope with different kinds of uncertainty, researchers developed several static and run-time analyses that quantify the level of noise, reliability, or accuracy. We survey the existing techniques in Section 7. These existing techniques suffer from one or more

of the following problems: 1) they have been developed *only for sequential* programs, 2) they are either *imprecise* (static analyses) or *lack guarantees* on result quality and soundness of monitoring code (empirical analyses), or 3) their applicability is *limited* – a single analysis is defined exclusively for a *specific* source of uncertainty (e.g., an unreliable instruction or a noisy sensor) and cannot be combined with others. Directly extending and generalizing the existing frameworks to a distributed setting can lead to subtle problems and/or run-time inefficiencies. An intriguing question is how to design a general analysis framework that will overcome these challenges, thus enabling a flexible and precise uncertainty analysis for parallel computations.

Our Work. We present Diamont, the first system for *sound, precise and efficient* runtime monitoring of uncertainty in *distributed applications*. Diamont offers a flexible runtime system for specifying and verifying uncertainty bounds in the face of *various sources of uncertainty*. Diamont supports programs consisting of distributed processes that communicate via asynchronous message-passing. Each process communicates with the others using strongly-typed communication channels through the common **send** and **receive** communication primitives. Diamont includes multiple language constructs for dynamic monitoring:

- **Dynamic types and data channels:** The developer specifies the variables that need to be dynamically monitored by annotating them using the **dynamic** type qualifier. In addition, Diamont introduces dynamic channels that use specialized communication primitives to reliably transfer the monitoring information.
- **Runtime Monitoring of Uncertainty:** Diamont maintains *uncertain intervals* for dynamically monitored variables – these map variables to a maximum error bound and a probability that the error is within the bound. Diamont propagates this uncertainty through computations. It can precisely do so even for individual array elements and unbounded loops – factors that usually reduce precision of existing analyses like Parallely [19] and DECAF [6].
- **Checkers:** Diamont’s **check** statement evaluates logical predicates over the program state and the monitored uncertainty to report violations. For example, the check can verify whether the magnitude of a variable’s error is less than a developer-defined threshold. Using Diamont’s checks, developers can decide if further attention should be given to the results. If the uncertainty of a result is acceptable at runtime, developers can avoid costly error checking and correction mechanisms.

We implemented Diamont for a distributed fragment of the Go language, extended with the dynamic type and check statements. Diamont performs static analysis at the level of an intermediate representation (IR) extracted from the Go code. It generates instrumented Go code with dynamic monitoring implemented via a Go library.

Diamont also presents a set of optimizations to reduce the runtime overhead arising from the monitoring of uncertain intervals throughout and across processes. These optimizations include: 1) combining static analysis with dynamic monitoring 2) approximating dynamically monitored uncertainty of arrays, 3) moving check statements across processes, and 4) using compiler techniques such as constant propagation and dead-code elimination. These optimizations give Diamont a significant advantage over direct extensions of systems like Decaf [6] or AffineFloat [13] to

parallel programs. However, developers who try to manually implement such run-time system optimizations that span multiple processes can easily make subtle errors.

Verified Runtime and Optimizations. We prove the soundness of the Diamont runtime and optimizations. Soundness of a Diamont program means that if the execution passes a variable uncertainty check, then the uncertainty of the variable is within the bound specified in the check statement. An optimization is sound if all check failures in a program are also guaranteed to occur in its optimized version.

Diamont’s runtime system is sound for programs that satisfy the *symmetric non-determinism* property [3] – i.e., each receive statement must have a unique matching send statement, or a set of symmetric matching send statements. Many common parallel patterns in data analytics applications [19, 34] satisfy this property. We use *canonical sequentialization* [3, 19], which rewrites a symmetrically nondeterministic parallel program to an equivalent sequential program. We can then prove soundness of runtime monitoring on the sequentialized program. Lastly, we show that this soundness proof also applies to the original parallel program.

Through sequentialization, Diamont can also automatically verify type safety and the absence of deadlocks of programs caused by approximations, the runtime system, or optimizations that change communication patterns.

Results. We applied Diamont on eight parallel applications. These real-world applications come from the domains of graph analytics, precision agriculture, and media processing. We modeled four sources of uncertainty: noisy communication, precision reduction (compression), noisy inputs, and timing errors.

We showed that Diamont can verify important end-to-end properties for all applications. In particular, we looked at four error probability predicates of end results, three error magnitude predicates, and one predicate on both error probability and magnitude. These properties cannot be validated by existing static techniques [10, 19, 26].

Our optimizations reduced the runtime overhead of Diamont with respect to the unmonitored program. Directly extending existing sequential runtime analyses to parallel settings leads to overheads between 30-80%. Our optimizations reduced the overhead to a geomean of 3% and maximum of 16.3% while satisfying strict predicates. We show that these overheads remain low and the communication of monitoring data is minimized even when the input size increases, especially for applications that implement intensive communication. These results demonstrate that even in the face of both uncertainty and significant parallelism, runtime monitoring is still practical.

Contributions. The paper makes several contributions:

- **Diamont.** Diamont is a system for dynamically monitoring uncertainty properties in strongly-typed, message-passing, asynchronous programs. We show that Diamont can soundly monitor uncertainty (error probability and magnitude).
- **Optimizations for reducing overhead.** We present several optimizations that reduce the overhead of performing runtime monitoring across processes.
- **Implementation.** We implement Diamont’s analysis and runtime system with optimizations for a subset of Go.
- **Evaluation.** We evaluate Diamont on 8 benchmarks. We show that Diamont can verify important correctness properties with small runtime overheads.

```

1  var Q [NUMSENSORS] process; var R [NUMWORKERS] process
2  type point struct { /*@dynamic*/ temperature, humidity float64 }
3
4  func Manager { // declarations & setup skipped to preserve space
5      for i, IoTDevice := range(Q) { data[i] = receive(IoTDevice) }
6      centers = // randomly pick some nodes
7      for i, Worker := range(R) { send(Worker, data) }
8      for j:=0; j<ITERATIONS; j++ {
9          for _, Worker := range(R) { send(Worker, centers) }
10         for i, Worker := range(R) { newcenters[i] = receive(Worker) }
11         centers = AverageOverThreads(newcenters)
12     }
13     checkArr(centers, 1, 0.99, 4, 0.99)
14 }
15
16 func IoTDevice {
17     /*@dynamic*/ var temperature, humidity float64
18     tempVal, tempErr, tempConf := readTemperature()
19     humidVal, humidErr, humidConf := readHumidity()
20     temperature = track(tempVal, tempErr, tempConf)
21     humidity = track(humidVal, humidErr, humidConf)
22     send(Manager, point{temperature, humidity})
23 }
24
25 func Worker {
26     var data [NUMSENSORS] point
27     var centers, newcenters [NUMCENTERS] point
28     /*@dynamic*/ var assign [PERTHREAD] int
29     data = receive(Manager)
30     for iter:=0; iter<ITERATIONS; iter++ {
31         centers = receive(Manager)
32         newcenters = kmeansKernel(data, centers, assign)
33         send(Manager, newcenters)
34     } }

```

Fig. 1: K-means Algorithm in a Smart Agriculture Setup in the Go Language

2 Example

We consider a scenario from precision agriculture [20]. Multiple low-power embedded systems with sensors are distributed across a field to monitor changes in the environment. Each embedded system (e.g., Raspberry Pis) can read the temperature, humidity, or other properties using their sensors. It can perform limited local processing of the readings, and periodically sends those results to a server for further (typically more expensive) analysis.

Figure 1 shows an implementation of the application in Go. The program has multiple parallel processes that communicate over typed channels using the Diamond API using matched `send` and `receive` statements (E.g., Lines 5, 22). The `Manager` process coordinates the computation.

The process group `Q` is of a set of processes running on embedded systems `IoTDevice1, ..., NUMSENSORS` that read sensor values and communicate the data to the `Manager`. Each `IoTDevice` gathers and stores datapoints using the struct `point`

from Line 2. The `/*dynamic*/` annotation indicates that the fields of `point` are of `dynamic` type. Diamont monitors the uncertainty of `dynamic` variables at runtime.

The `Manager` process first gathers sensor data (Line 5) from each `IoTDevice`. Then it performs a distributed k-means clustering analysis using the processes in the group `R`. The `Manager` picks a set of random points as the initial cluster centers (Line 6). Next, over `ITERATIONS` iterations, it updates the cluster centers (Lines 8-12).

Each `Worker` process from the group `R` processes a subset of the data points to calculate new cluster centers (Lines 30-33) for that subset. The `Manager` combines the partial results from each `Worker` and redistributes them (Line 11).

2.1 Sources of Uncertainty

Approximate sensors. Sensors are often noisy (e.g., the AM2302-DH22 relative humidity and temperature sensor has an error range of $\pm 0.5^\circ\text{F}$ for temperature and $\pm 2\% \text{RH}$ for humidity reading [24]). Each process in `Q` calculates the error of its sensors while reading the value at Lines 18 and 19. This error calculation can come from the sensor specification (e.g. [24]). Next, Lines 20 and 21 initialize `dynamic` variables using the sensor value and error.

Approximate Communication. We also consider the impact of communication over noisy channels (Line 7, 29), prevalent in situations where sensors are deployed in remote areas (E.g., [45]). Messages in such channels can be corrupted with a small probability [29]. Instead of implementing costly error correction mechanisms, a developer may choose to deal with potentially incorrect data to save resources.

An uncertainty model ψ provides parameters such as the probability of message corruption. For example, $\psi(\text{Manager}, \text{Worker}, \text{dynamic float}\langle 64 \rangle) = 1 - 10^{-7}$ indicates that the probability of corruption of a `dynamic float<64>` type message from `Manager` to `Worker` is 10^{-7} . The specification is modeled after the ones from [5, 10, 37].

2.2 Verification

Properties. We wish to verify that the final values of `centers` are close to the true cluster centers with high probability. We encode this requirement in the `checkArr` statement in Line 13. This check specifies a maximum error magnitude and probability for each `dynamic` field in the struct. This program has features that make static verification using tools such as Parallely [19] challenging:

- The error specification of the sensors may not be known a priori. Additionally, prior static verification techniques require worst-case bounds for the number of loop iterations and the number of processes. Using worst-case estimates for these in a static analysis will invalidate many correct programs.
- Parallely treats entire arrays as single variables, and thus array analysis accumulates errors even across two different array locations. Consequently, the conservative static estimate of uncertain intervals quickly expands to unusable levels for any sufficiently large number of sensors for our example.




Fig. 2: Diamont Workflow

Workflow. Diamont combines static and dynamic analyses to verify safety and accuracy properties at *runtime*. Figure 2 shows the workflow for generating an instrumented program in Diamont. Given a Go program, Diamont 1) translates it to Diamont-IR, 2) sequentializes the program to statically verify type safety, deadlock-freeness, and the applicability of the runtime analysis, and 3) produces an instrumented version of the original Go program with an *uncertainty map* for each process. The sequentialized version of the code in Figure 1 is in Appendix F [18].

The *uncertainty map* of a process maintains a conservative uncertain interval for each **dynamic** local variable. Uncertain intervals are stored as pairs $\langle d, r \rangle$ indicating that the maximum error of the associated variable is $\leq d$ with probability $\geq r$. The default uncertain interval is $\langle 0, 1 \rangle$ (no error with 100% confidence). Developers can use `track` statements (E.g., Line 20) to use external error specifications within Diamont. When a dynamic variable is updated, Diamont also updates the uncertain interval. Diamont’s instrumentation 1) initializes the uncertain interval of the data in `IoTDevice`, 2) communicates the uncertain interval across process boundaries, 3) propagates this uncertainty through computations, and 4) checks the uncertain interval of the array at the end of the program against a developer-specified bound.

We verified this system for a setting with 128 sensors and a set of 8 workers performing the k-means computation over 10 iterations. As more and more computations containing unreliable values affect the `centers` array, the uncertain interval of individual elements widens. However, the specification is still satisfied.

Overhead. Diamont’s instrumentation adds runtime overhead. To reduce overhead, Diamont applies optimizations such as constant propagation, dead code elimination, and simplification of monitoring uncertainty in arrays. To reduce overhead when transmitting arrays, Diamont transmits the maximum uncertainty among the elements of the array as the uncertainty of every element of the array. This allows Diamont to only communicate one uncertain interval across processes, while maintaining high analysis precision in other parts of the program. These optimizations reduce Diamont’s overhead from 42% to 3.2%. Increasing the number of sensors does not significantly increase overhead (Section 6.3). Even for 2-8x larger data, the overhead remains below 5%.

3 Diamont System

Diamont takes as input a Go program and an uncertainty model. Diamont first converts the program to the Diamont-IR and verifies important safety properties necessary to ensure that the runtime system will be sound. Finally, Diamont generates instrumented Go code. The full syntax and semantics of Diamont are available in Appendix A [18].

$m, v \in \text{NUFU}\{\emptyset\}$	<i>values</i>	$S \rightarrow T \ x \mid T \ a[n^+]$	<i>declarations</i>
$Exp \rightarrow m \mid \langle m, v \rangle \mid x$	<i>expressions</i>	$\mid x = Exp$	<i>assignment</i>
$\mid Exp \ op \ Exp$		$\mid \text{dyn-send}(\alpha, T, x)$	<i>send dynamic</i>
$AEx \rightarrow d \mid d \cdot x \mid d \cdot a[Exp^+]$	<i>affine</i>	$\mid x = \text{dyn-recv}(\alpha, T)$	<i>receive dynamic</i>
$\mid AEx \pm AEx$	<i>expressions</i>	$\mid x = \text{rdDyn}(y)$	<i>read dynamic map</i>
$q \rightarrow \text{precise} \mid \text{approx}$	<i>type</i>	$\mid x = \text{endorse}(y)$	<i>cast to precise</i>
$\mid \text{dynamic}$	<i>qualifiers</i>	$\mid x = \text{track}(y, \langle d, r \rangle^+)$	<i>initiate monitoring</i>
$t \rightarrow \text{int}\langle n \rangle \mid \text{float}\langle n \rangle$	<i>basic types</i>	$\mid x = Exp? \ Exp : Exp$	<i>conditional choice</i>
$T \rightarrow q \ t \mid q \ t \ \square \mid \text{struct } T^+$	<i>types</i>	$\mid \text{check}(AEx, \langle d, r \rangle^+)$	<i>check error</i>
$P \rightarrow [S]_\alpha$	<i>process</i>	$\mid \text{checkArr}(a, \langle d, r \rangle^+)$	<i>check array error</i>
$\mid \Pi. \alpha : X \ [S]_\alpha$	<i>process group</i>		
$\mid P \parallel P$	<i>parallel comp</i>		

Fig. 3: Diamont-IR Syntax Extensions (full language contains conditionals, loops and function calls)

3.1 Syntax

Go Language. Diamont supports a subset of the Go Programming Language (matching the features of Diamont-IR along with external functions that do not perform communication) extended with an API for distributed communication and annotations in comments for type qualifiers.

Diamont-IR. Diamont’s intermediate representation supports a strongly typed imperative language with primitives for asynchronous communication. Diamont extends the syntax of Parallely [19] with support for the additional **dynamic** type. Figure 3 defines the subset of Diamont syntax dealing with **dynamic** data. Here, d refers to reals, r to probabilities, n to positive integers, x, y to variables, and a to array variables. The full syntax includes conditionals, loops, operations on arrays, and structs.

Types. Diamont’s type qualifiers explicitly split data into either **precise** (no uncertainty), **dynamic** (uncertainty monitored at runtime), or **approx** (uncertain but unmonitored). Diamont’s type system ensures that uncertainties in executions do not cause errors in critical program sections and ensures that the dynamic monitoring is sound by avoiding control flow divergence. Using type inference, Diamont automatically annotates some variables as dynamic to reduce programmer burden.

Communication. Processes communicate by sending and receiving messages over typed channels. For each pair of processes, Diamont provides a set of logical sub-channels for communication, further split by message type (μ). A **send** statement asynchronously sends a value to another process using a unique process identifier. The receiving process uses the blocking **receive** statement to read the message. Diamont supports communication of **dynamic** type data through **dyn-send** and **dyn-recv** statements, which also send the monitored uncertainty using reliable channels.

Type conversion. To explicitly convert a variable to **dynamic** type, the developer or compiler can use a **track** statement ($x = \text{track}(y, \langle d, r \rangle)$), which sets the uncertain interval to $\langle d, r \rangle$. **track** statements can be used to initiate monitoring for variables updated by external functions, or to incorporate informal specifications (e.g., from a datasheet) into Diamont. Similarly, the **endorse** statement ($x = \text{endorse}(y)$) converts an **approx** or **dynamic** variable to a **precise** variable, usually after a user-defined

$$\begin{array}{c}
\text{S-ASSIGN-DYN} \\
\frac{(x, \dots) \in D \quad \langle e, \sigma, h \rangle \Downarrow v \quad d = \langle \text{calc-eps}(e, D), \text{calc-del}(e, D) \rangle \quad D' = D[x \mapsto d] \quad \langle n_b, \langle 1 \rangle \rangle = \sigma(x) \quad h' = h[n_b \mapsto v]}{\langle x = e, \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{\psi} \langle \text{skip}, \langle \sigma, h' \rangle, \mu, D' \rangle} \\
\\
\text{S-DYNRECEIVE} \\
\frac{\mu[\langle \beta, \alpha, D_t \rangle] = d :: m_d \quad \mu' = \mu[\langle \beta, \alpha, D_t \rangle \mapsto m_d] \quad d_b = \langle d, \epsilon, d, \delta \times \psi(\beta, \alpha, t) \rangle \quad D' = D[x \mapsto d_b]}{\langle [x = \text{dyn-recv}(\beta, t)]_\alpha, \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{\psi} \langle [x = \text{receive}(\beta, t)]_\alpha, \langle \sigma, h \rangle, \mu', D' \rangle} \\
\\
\text{S-CHECK-FAIL} \\
\frac{\text{calc-eps}(AEx, D) > d \vee \text{calc-del}(AEx, D) < r \quad \langle \text{check}(AEx, d, r), \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{\psi} \langle \text{skip}, \perp, \mu, D \rangle} \\
\\
\text{S-CAST} \\
\frac{\langle n_b, \langle 1 \rangle \rangle = \sigma(y) \quad h[n'_b] = m \quad m' = \text{cast}(T, m) \quad \langle n_b, \langle 1 \rangle \rangle = \sigma(x) \quad h' = h[n_b \mapsto m'] \quad d = \langle \text{cast-eps}(x, y, D), D[y].\delta \rangle \quad D' = D[x \mapsto d]}{\langle x = (\text{dynamic } T)y, \langle \sigma, h' \rangle, \mu, D' \rangle \xrightarrow{\psi} \langle \text{skip}, \langle \sigma, h' \rangle, \mu, D' \rangle}
\end{array}$$

Fig. 4: Semantics of Dynamic Monitoring (Selection)

check (similar to EnerJ [37]). The `rdDyn` intrinsic (`rdDyn(x)`) can be used to read the monitored uncertainty of a dynamic variable.

Uncertainty Model (ψ). It specifies the reliability/accuracy of program components (e.g., the probability of message corruption or the probability that a sensor fails).

Specifications. Diamont exposes the following statements to check specifications of dynamically monitored variables.

- `check(AEx, ⟨d, r⟩)`: It checks if an affine expression AEx has a maximum error $\leq d$ with probability $\geq r$. If the specification is not satisfied, the check fails.
- `checkArr(a, ⟨d, r⟩)`: It checks if the dynamically monitored uncertainty for *each* element in array a satisfies the specification.

While this version of Diamont stops the execution if a check fails, it can be extended to trigger a recovery mechanism instead [1, 15, 22]. Aloe [22] represents recoverable computations with blocks of the form `try { ... } check (...) recover { ... }`. Using this construct, Diamont can recover the execution if a check fails, and calculate the effect of (possibly imperfect) checks and recovery mechanisms on uncertainty. Formalization of recovery for distributed programs, however, is out of scope of this paper.

Structs. The programmer can specify the uncertainty of each field of a struct in a `track` statement by using multiple $\langle d, r \rangle$ pairs. The programmer can check each field of a struct in `check` and `checkArr` statements in a similar manner.

3.2 Diamont Semantics

Semantics for `precise` and `approx` data in Diamont are the same as those from Parallely [19]. For `dynamic` data, the compiler adds instructions to monitor their uncertain intervals alongside the original program instructions.

References, Frames, Stacks, and Heaps. A *reference* is a pair $\langle n_b, \langle n_1, \dots, n_k \rangle \rangle \in \text{Ref}$ that contains a base address $n_b \in \text{Loc}$ and dimension descriptor $\langle n_1, \dots, n_k \rangle$ denoting the location and dimension of variables in the heap. A *frame* $\sigma \in \mathbb{E} = \text{Var} \rightarrow \text{Ref}$

$$\text{calc-eps}(e,D) = \begin{cases} 0 & e \text{ is a constant} \\ D[x].\epsilon & e \text{ is a variable } x \\ D[x].\epsilon + D[y].\epsilon & e \text{ is } x \pm y \\ |x| \times D[y].\epsilon + |y| \times D[x].\epsilon + D[x].\epsilon \times D[y].\epsilon & e \text{ is } x \times y \\ \infty & e \text{ is } x \div y \wedge 0 \in [y \pm D[y].\epsilon] \\ \frac{(|x| \times D[y].\epsilon + |y| \times D[x].\epsilon)}{(|y| \times (|y| - D[y].\epsilon))} & e \text{ is } x \div y \wedge 0 \notin [y \pm D[y].\epsilon] \end{cases}$$

$$\text{calc-del}(e,D) = \max(0, (\sum_{x \in \rho(e)} D[x].\delta) - (|\rho(e)| - 1))$$

$$\text{cast-eps}(x,v,D) = \max(\max(x + D[x].\epsilon, v + D[x].\epsilon) - v, v - \min(x - D[x].\epsilon, v - D[x].\epsilon))$$

Fig. 5: Runtime for Dynamic Monitoring of Uncertainty

maps program variables to references. A *heap* $h \in H = \mathbb{N} \rightarrow \mathbb{N} \cup \mathbb{F} \cup \{\emptyset\}$ is a finite map from addresses to values (Integers, Floats or the special *empty message* $[\emptyset]$). Each process i maintains its own private environment consisting of a frame and a heap $\langle \sigma^i, h^i \rangle \in \Lambda = \{H \times E\} \cup \perp$, where \perp is considered to be an error state.

Uncertainty Map. For each process, Diamont defines an *uncertainty map* (D) to attach each variable with a uncertain interval, consisting of a maximum absolute error (ϵ), and a probability/confidence (δ) that the true error is below ϵ .

Local Semantics. The small-step relation $\langle s, \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{p}_\psi \langle s', \langle \sigma', h' \rangle, \mu', D' \rangle$ defines a process in the program evaluating in its local frame σ , heap h , uncertainty map D , and the global channel set μ . Figure 4 presents a selection of the semantics.

- **Initialization:** Each **dynamic** variable is initialized by setting the maximum error ϵ to 0 and the confidence δ to 1.
- **Expressions:** The **S-Assign-Dyn** rule in Figure 4 is applied when a **dynamic** variable is updated by assigning it an expression e . We use a big-step evaluation relation of the form $\langle e, \sigma, h \rangle \Downarrow v$ to compute the result of the expression. Diamont supports typical integer and floating point operations. For **dynamic** variables, in addition to the assigned variable, Diamont updates its interval using the uncertain interval arithmetic defined in Figure 5. The **calc-eps** function is used to calculate an expression's maximum error. The confidence in this maximum error is then computed using **calc-del** ($\rho(e)$ returns the list of variables used in an expression e .) To avoid any assumptions about the independence of the uncertainties (prior approaches such as [6] restrictively assumed all the operations and probability of failures are independent) Diamont uses the conservative union bound.
- **Communication:** When sending **dynamic** variables of type T to another process (rule **S-DynSend**), Diamont uses special channels (D_T) that are assumed to be fully reliable to communicate the relevant uncertain intervals before sending the data.¹ At the receiver (rule **S-DynReceive**), Diamont updates the local uncertainty map. Diamont assumes the channel failure rate is independent of the message content and reduces the confidence based on the failure rate defined in the Uncertainty Model.

¹ ++ denotes adding a element to the end of the message queue.

- **Precision Manipulation:** Diamont monitors the errors introduced to programs through *cast* statements that change the precision of values of the same general type (int or float). In the rule **S-Cast**, the added error is calculated using the `cast-eps(x, v, D)` function using the casted value v and the original variable x . Confidence remains the same.
- **Conditionals:** For branching on **dynamic** values, Diamont supports an operator $x = \text{cond? } e_1 : e_2$ (conditional choice) where `cond` compares a **dynamic** value against a threshold. We check if the *entire* interval associated with the value is greater or less than the threshold. If neither case is true, we compute both expressions and the interval of x becomes the smallest closed interval that contains all possible intervals.
- **Checks:** If a check fails, the Diamont program transitions into an error state (Figure 4 rule **S-Check-Fail**). To prevent such check failures, the user can implement error recovery mechanisms.

Global Semantics. We define a global configuration as $\langle \epsilon, \mu, \omega, P \rangle$, consisting of a global environment $\epsilon \in Env = Pid \mapsto A$, a set of typed channels $\mu \in Channel = Pid \times Pid \times Type \rightarrow Val^*$, global uncertainty map $\omega \in Pid \mapsto D$, and the program P . Small step transitions of the form $\boxed{(\epsilon, \omega, \mu, P) \xrightarrow{\alpha, r}_{\psi} (\epsilon', \omega', \mu', P')}$ define a process α taking a step and thus changing the global configuration. Inter-process communication happens using the typed channels – though processes adding to and reading from the relevant queue. Complete semantics are available in Appendix A.

3.3 Soundness of Runtime Monitoring

Diamont’s runtime system works across distributed processes. We use *Canonical Sequentialization* [3] to simplify our reasoning about the soundness of the runtime system. Canonical sequentialization uses the assumption that correct programs tend to be well-structured to generate a sequential program that over-approximates the semantics of a parallel program. If such a sequentialized program can be generated, then the parallel program is deadlock-free, and local safety properties that hold for the sequentialized program also hold for the parallel program.

To be sequentializable, the parallel program must be *symmetrically nondeterministic* – each receive statement must only have a single matching send statement, or a set of symmetric matching send statements². We use a set of rewrite rules of the form $\boxed{\Gamma, \mathcal{S}, P \rightsquigarrow \Gamma', \mathcal{S}', P'}$ to rewrite a parallel program P to a sequential program \mathcal{S}' step by step (the rules are available in Appendix C). The *context* Γ is used as a symbolic set of messages in flight, and P' is the part of the parallel program that remains to be rewritten. The sequentialization process applies the rewrite steps until the entire program is rewritten to \mathcal{S}' . We extend the results from prior work [3, 19] to show that rewrite rules maintain equivalent behavior between the original parallel program and the generated sequential program, i.e., they both produce the same environment and uncertainty map at the halting states of the programs.

² Many popular parallel application patterns (e.g. Map, Reduce, Scatter-Gather, Stencil) exhibit symmetric non-determinism [3, 19]. Further, programs satisfying this property can be less error-prone [3].

$$\begin{array}{l}
S = [] \\
P = \left[\begin{array}{l} \text{int } \alpha.n = 1 \text{ [r] } 0; \\ \text{send}(\beta, \text{int}, \alpha.n); \end{array} \right]_{\alpha} \parallel \left[\begin{array}{l} \text{int } \beta.x; \\ \beta.x = \text{receive}(\alpha, \text{int}); \end{array} \right]_{\beta} \rightsquigarrow^* \begin{array}{l} S = \left[\begin{array}{l} \text{int } \alpha.n = 1 \text{ [r] } 0; \\ \text{int } \beta.x; \\ \beta.x = \alpha.n; \end{array} \right] \\ P = [\text{skip};] \end{array}
\end{array}$$

Fig. 6: Canonical Sequentialization: An Example of the Rewriting Process.

Figure 6 shows a small program with inter-process communication (P) and its canonical sequentialization (\mathcal{S}) generated using the rewrite rules. We show that the existence of a canonical sequentialization guarantees that uncertain intervals are not affected by the different possible interleavings of processes during execution, allowing us to generate correct monitoring code.

In contrast, consider the following program where the process α has a receive statement that receives from two other processes:

$$\left[\alpha.\text{res} = \text{receive}(*); \right]_{\alpha} \parallel \left[\begin{array}{l} \beta.\text{out} = \text{func1}(); \\ \text{send}(\alpha, \beta.\text{out}); \end{array} \right]_{\beta} \parallel \left[\begin{array}{l} \gamma.\text{out} = \text{func2}(); \\ \text{send}(\alpha, \gamma.\text{out}); \end{array} \right]_{\gamma}$$

The final value of `res` depends on the runtime interleavings and it is difficult to generate monitoring code at compilation time that soundly calculates an uncertain interval combining all possible interleavings. Therefore, we limit our analysis only to programs with canonical sequentializations and prove that the runtime is sound.

We use the notation developed in Chisel [26] to state the following soundness theorem. Recall that Diamont’s runtime monitors two properties for each `dynamic` variable x : (1) the maximum possible error magnitude ($D[x].\epsilon$) and (2) a probability ($D[x].\delta$) that the *precise* value of x is within $x \pm D[x].\epsilon$. The notation $\Delta(x)$ denotes the *true error* of a variable x , and $\llbracket \mathcal{R}^*[E] \rrbracket(\sigma, \varphi)$ denotes the *true probability* that an environment σ sampled from the *environment distribution* φ satisfies the error comparison E .

Theorem 1 (Soundness of dynamic monitoring). *For programs not containing `track` and `endorse` statements, for all statements s , and for all x s.t. $\Theta \vdash x$: dynamic t , $\Theta \vdash s$: Θ' and $\langle s, \langle \sigma, D, \varphi \rangle \rangle \Downarrow \langle s', \langle \sigma', D', \varphi' \rangle \rangle \implies \llbracket \mathcal{R}^*[D'[x].\epsilon \geq \Delta(x)] \rrbracket(\sigma', \varphi') \geq D'[x].\delta$*

First, we use induction over the sequential subset of Diamont to show that, if the program s type checks, and evaluates in the global environment σ and uncertainty map D to s' , resulting in the environment σ' and uncertainty map D' , then, for all dynamic variables x , the *true error* of x is at most by $D'[x].\epsilon$ with probability at least $D'[x].\delta$. This indicates that we soundly over-approximate the uncertainty of x .

Next, we utilize canonical sequentialization to prove that the theorem holds for the parallel subset of the language as well. First, we extend the results from [19] to prove that if we can rewrite a parallel program P into a sequential program \mathcal{S} , then P and \mathcal{S} have equivalent behavior. We use this fact to reason that our proof of soundness for the sequential subset of Diamont is also applicable to parallel programs that can be canonically sequentialized. Therefore, Theorem 1 holds and our overall analysis is sound (full proof is available in Appendix D).

Our analysis only applies to programs with `track` and `endorse` statements if developers use them in a sound manner. For `track` statements, developers must




Fig. 7: Optimizations Using Static Analysis in Diamont.

ensure that the bounds they provide are a sound over-approximation of the true uncertainty at that program point. As in prior work [37], by inserting `endorse` statements, developers certify that treating the relevant `approx` or `dynamic` value as `precise` is always safe and will not result in undesirable behavior.

4 Optimizations for Reducing Overhead

We implemented several optimizations that transform the programs to reduce the overhead of dynamic monitoring and proved them to be sound.

Communication. When communicating large `dynamic` type arrays, Diamont must also communicate the uncertain interval for each array element, resulting in a large communication overhead. One way to reduce this overhead is to calculate a single conservative approximation of the set of uncertain intervals for the array elements. For example, the maximum error of any element of an array can be soundly over-approximated by the largest maximum error among all of its elements (similarly, the smallest error confidence). The process sending the data calculates the conservative approximation while using the regular communication primitives for the data. At the end it sends the conservatively approximate uncertain interval. At the receiver, this uncertain interval is taken as the uncertain interval of *each* element in the received array and the compiler adds `track` statements to restart dynamic monitoring.

This optimization does not approximate the uncertain interval of the array at all program points, rather it affects only communication statements. Even with the resulting loss in precision of the analysis, Diamont still achieves better results than existing static analyses which use a single uncertain interval for arrays through the *entire* program.

Utilizing static analysis. We can further reduce overheads by exploiting common communication patterns. For example, the program at the top of Figure 7 contains a remote procedure call. Process α sends an input to process β , which applies the function `fn` to the input and returns the value. Transferring uncertain intervals along with the data can become expensive if many such calls are made.

We use existing static analysis techniques [10, 19, 26] to analyze only the remote function call and generate function specifications (precise semantics are in Appendix A Figure 9), even if they are unable to analyze the entire program. Consider the transformed program at the bottom of Figure 7. Using the specification, Diamont produces the same behavior as the original program by generating code to 1) check if the specification requirements are satisfied (Line 8), 2) transfer the data as `approx`

type (Line 9), 3) compute without dynamic monitoring, and 4) re-initialize dynamic monitoring using the error guarantees from the specification (Line 11).

This optimization can be safely used when the function performs no communication and has no other side effects. However, it may not be possible to verify some static specifications at runtime. For example: the runtime will not be able to calculate $\mathcal{R}^*[d_i \geq \Delta(\text{dat})]$ for some values for d_i . Therefore, this optimization may introduce some imprecision to the dynamic monitoring.

Early checking. For a subset of instructions we can perform static analysis to stop runtime monitoring earlier. We perform this task by *moving up* the check to the earliest possible location using a set of rewrites. This rewrite rule is one such example:

$$\left[\begin{array}{l} \alpha.x = \alpha.a + \alpha.b; \\ \text{check}(\text{AExp}, d, r); \end{array} \right] \Rightarrow \left[\begin{array}{l} \text{check}(\text{AExp}[(\alpha.a + \alpha.b)/\alpha.x], d, r); \\ \alpha.x = \alpha.a + \alpha.b; \end{array} \right]$$

In this rule, Diamont looks for a check immediately following an addition. Since the error magnitude of the result of the addition is the sum of the error magnitudes of the variables that are being added, we can substitute the result variable $\alpha.x$ in the check with $\alpha.a + \alpha.b$. As the `calc-del` function of the runtime looks for the set of variables in the specification (AExp), the error probability is calculated correctly as well. Diamont can now safely move the check before the addition

These re-write rules closely follow the static analysis as defined and proven sound in [19] for the sequential subset of the language (Appendix D.2.2). This optimization reduces updates to the uncertainty map as monitoring can be stopped after the check is performed. However, it can only be applied when the check refers to variables from a single process. Further, the check cannot be moved up if error calculations depend on the value of variables (as in multiplication/division).

Debloating and compiler optimizations. Diamont further reduces overhead by using constant propagation and dead code elimination to remove unnecessary updates to the uncertainty map. In addition, Diamont eliminates either error magnitude monitoring or confidence monitoring based on the checks in the program. For example, if all checks require the error magnitude to be zero (reliability in [10]) Diamont will only calculate confidence at runtime.

4.1 Soundness

For each optimization we show that both the original program (s) and the optimized version (s_{opt}) produce the same behavior, i.e., if the original program fails a check, the optimized version is also guaranteed to fail. Canonical sequentialization makes such proofs easier. Formally, we define the soundness of an optimization as follows:

Definition 1 (Optimization soundness). *For a program s and its optimized version s_{opt} , $\langle s, \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{*}_{\psi} \langle s', \perp, \neg, \neg \rangle \implies \langle s_{opt}, \langle \sigma, h \rangle, \mu, D \rangle \xrightarrow{*}_{\psi} \langle s'', \perp, \neg, \neg \rangle$*

This definition states that if there is an execution where the original program s starting from an environment σ , heap h , uncertainty map D , and the global channel set μ evaluates to s' and enters into the error state (\perp), the optimized version s_{opt}

$$s^{seq} = \left[\begin{array}{l} \beta.\text{dat} = \alpha.\text{in}; \\ \beta.\text{res} = \text{fn}(\beta.\text{dat}); \\ \alpha.\text{out} = \beta.\text{res}; \\ \text{check}(\alpha.\text{out}, d_{\text{check}}, r_{\text{check}}); \end{array} \right] \quad s_{opt}^{seq} = \left[\begin{array}{l} \text{check}(\alpha.\text{in}, d_i, 0); \\ \beta.\text{dat} = \alpha.\text{in}; \\ \beta.\text{res} = \text{fn}(\beta.\text{dat}); \\ \alpha.\text{tmp} = \beta.\text{res}; \\ \alpha.\text{out} = \text{track}(\alpha.\text{tmp}, d, r * \text{rdDyn}(\alpha.\text{in}).\delta); \\ \text{check}(\alpha.\text{out}, d_{\text{check}}, r_{\text{check}}); \end{array} \right]$$

Fig. 8: Example Sequentializations Used in the Proofs

starting from the same state σ , heap h , and D must also enter the error state (even if the final channel or uncertainty map states differ).

For each optimization, we show that the pairs s and s_{opt} are sound according to this definition. Consider the static analysis based optimization in Figure 7. Proving the soundness of this optimization requires us to show that the two parallel programs produce the same result with regards to the dynamic monitoring. We can simplify this process significantly by using sequentialization. We first show that the two versions of the program can be sequentialized to s^{seq} and s_{opt}^{seq} in Figure 8. These sequentializations produce final environments that are equivalent to the original versions as proven in Lemma 1 (Appendix C). We can now simplify the proof to reasoning over the two sequential programs s^{seq} and s_{opt}^{seq} . We can next argue over all executions resulting in a check failure in s^{seq} and show that they result in a check failure in s_{opt}^{seq} (The full proofs are in Appendix D).

5 Methodology

Implementation and Testing Setup We parsed and translated Go programs written using a library of Diamont primitives to Diamont-IR using ANTLR. We used Python to sequentialize Diamont programs for checking properties such as type safety and deadlock-freedom, and then for generating instrumented Go code. We implemented distributed communication using RabbitMQ 3.8.7. We ran our experiments on a machine with a Xeon E5-1650 v4 CPU, 32 GB RAM, and Ubuntu 18.04. Each benchmark consisted of 8-10 worker processes.

Benchmarks We implemented a set of popular parallel benchmarks from prior literature that exhibit diverse parallel patterns and verified properties that quantify uncertainty in their executions (Table 1). We looked at the following benchmarks:

- *PageRank, SSSP, BFS*: Graph benchmarks commonly used in distributed Big Data applications. PageRank is used for search result optimization [27]. Single Source Shortest Path is used to make data routing decisions. Breadth First Search is used to find connected components in graphs. From CRONO [2].
- *SOR*: A kernel for successive over-relaxation. Used to extrapolate the state of a system over time. From Chisel [26].
- *Sobel*: Sobel edge-detection filter. From AxBench [44].
- *Matrix Mult.*: Multiplies two square matrices. Each worker process computes a subset of rows of the product.
- *Kmeans-Agri*: Partitions n-dimensional input points into k clusters (Section 2).
- *Regression*: Performs distributed linear regression on 2-D data. Each worker performs regression on a subset of data. The master thread averages the results.

Table 1: Benchmarks, Verified Properties, and Runtime Monitoring Overhead for Diamont. Baselines: \star :Decaf, \dagger :AffineFloat

Benchmark	Pattern	Verified Property	Overhead	
			Baseline	Diamont
PageRank	Scatter-Gather	checkArr(pagerank, 0, 0.9912)	30% \star	3.63%
SSSP	Scatter-Gather	checkArr(distance, 0, 0.9925)	33% \star	2.31%
BFS	Scatter-Gather	checkArr(visited, 0, 0.9925)	30% \star	4.06%
SOR	Stencil	checkArr(output, 1.19×10^{-7} , 1)	60% \dagger	3.49%
Sobel	Stencil	checkArr(output, 2.38×10^{-7} , 1)	71% \dagger	9.71%
Matrix Mult.	Map	checkArr(product, 6.6×10^{-6} , 1)	80% \dagger	16.27%
Kmeans-Agri	Map	checkArr(centers, (1.5, 0.9948), (2, 0.9948))	42% $\star\dagger$	3.32%
Regression	Map-Reduce	check(alpha, 0, 0.99) \wedge check(beta, 0, 0.99)	37% \star	0.45%

Inputs. The inputs for each benchmark used for our experiments are shown in Appendix E. For Section 6.3, we used larger inputs created by increasing the size of the array, the number of samples, or by using a larger input graph.

Sources of uncertainty. *Noisy channels* occasionally corrupt data sent over them (used for PageRank, SSSP, BFS, and Kmeans-Agri). We use a corruption rate of 10^{-7} . *Precision reduction* reduces floating point precision from 64-bit to 32-bit during communication only to save bandwidth (used in SOR, Sobel, Matrix Mult.). The *input* provided to the program itself can have inherent uncertainty. For Kmeans-Agri, we assume a 50:50 mixture of two different temperature-humidity sensors with different error specifications. *Timing errors* can cause the program to use stale or incomplete values (used for Regression).

Baselines. We compare the runtime of Diamont with optimizations to a baseline which is a straightforward parallel implementation of an existing static analysis via Diamont (either Decaf [6] or AffineFloat [13] without roundoff errors).

6 Evaluation

6.1 Can we verify important uncertainty properties using Diamont?

For each benchmark, we used Diamont to verify the properties shown in Column 3 of Table 1. Diamont successfully verified these properties on the final output of the program. Each check places an error magnitude and confidence bound on a single variable. For arrays each element must satisfy these bounds. For PageRank, SSSP, and BFS, the bounds ensure that key graph properties are calculated exactly $\geq 99\%$ of the time per node. For SOR, Sobel and Matrix Mult., the bounds limit the maximum error of the output due precision reduction. Kmeans-Agri was discussed in the example. For Regression, the bounds ensure that the output line parameters are correct $\geq 99\%$ of the time (high confidence is desirable for predictive models).

Parallely [19] cannot verify these properties. Diamont’s dynamic analysis of arrays and unbounded loops more effectively handles irregular input structure (e.g., graphs), which had to be conservatively bounded for static analysis. This allowed us to verify stronger properties for significantly bigger inputs than previously possible for existing reliability and accuracy static analyses. We observed that, even in the presence of errors, the error magnitude of the final outputs of our programs was acceptable.

Optimizations can affect the precision of the analysis. This effect is prominent in benchmarks with irregular computations (graph benchmarks). However, in our benchmarks, we found that baseline and optimized Diamont could verify nearly the same uncertainty bounds. For example, for BFS, Diamont could verify a confidence of 0.999 when using the baseline version. For benchmarks with regular computation patterns, such as SOR and Regression, there was no significant change.

In summary, *Diamont verifies important end-to-end uncertainty properties that cannot be verified using existing static analyses.*

6.2 What are the overheads associated with Diamont?

Columns 4 and 5 of Table 1 present the overhead of the baseline and optimized Diamont benchmarks respectively. Time for I/O and setup is excluded. Overhead is calculated as the percentage increase in runtime w.r.t. an unmonitored benchmark.

In our benchmarks, the runtime is dominated by communication, as is common in many distributed settings. In most cases, the runtime overhead for computing the uncertain intervals is a small fraction of the total runtime. Error magnitude calculation requires more computation than error confidence (see Figure 5). As a result, overhead for error magnitude benchmarks (SOR, Sobel, Matrix Mult.), is higher. This was especially true for the computationally intensive Matrix Mult.

Optimization impact. The Regression benchmark used a statically verified kernel error specification to eliminate monitoring. The communication optimization contributes around 98% of savings in all other benchmarks. Debloating also provided significant speedups. For example, without debloating PageRank is 3.9x slower and Sobel is 3.3x slower (our baseline is comparable to Diamont with debloating).

Are the overheads justified? Approximations have led to significant savings in prior work: 1) Communication: up to 62% performance improvement in approximate NoCs [11, 17], and 2) Computation: 2x speedup in loop perforation [40], 2.7x speedup in Paraprox [34], and up to 1.3x speedup from reduced precision in Precimonious [33]. As Diamont’s post-optimization overhead is lower than the speedups from these approximations, it can be used in conjunction with them to provide guarantees on the quality of results while still getting speedups.

In summary, *With optimization, overhead of Diamont analysis is at most 16.3% for our benchmarks, with a geomean of 3.04%.*

6.3 How does Diamont overhead depend on the program inputs?

Figure 9 shows the effect of input size on Diamont overhead. The X-Axis shows the relative input size and the Y-Axis shows overhead. The dashed and solid lines show the unoptimized baseline and optimized Diamont versions respectively. Each marker indicates a different benchmark. Overall, the overhead of the optimized versions is significantly lower than the baseline versions. Most optimized versions have an overhead less than 25% for all inputs. The table in Figure 9 shows the geomean of




Fig. 9: Input Size vs. Overhead. Table shows geomean overheads across programs.

the overhead across all benchmarks for different relative input sizes. While baseline overhead increases to an average of 94%, optimized overhead only reaches 14%.

For Matrix Mult., computation increases faster with input size than communication ($O(n^3)$ vs. $O(n^2)$). Thus the major source of overhead becomes the computation of the monitored uncertainty, rather than communication. This benchmark illustrates that Diamont is more useful in cases where the program is communication-bound.

The unoptimized baseline also sends significantly more data (3x to 5x) compared to the optimized version. This is due to the array communication optimization. The communication overhead of the optimized version is negligible.

In summary, *as input size grows, the improvement caused by optimizations on Diamont runtime performance increases over the baseline runtime system.*

7 Related Work

Several analyses are related (in part) to Diamont’s functionality, as shown in Table 2. Columns 2-4 indicate whether the analysis is static, empirical (sampling-based), or runtime based. Columns 5-6 indicate support for error confidence (reliability) and error magnitude (accuracy) analysis. Column 9 indicates if the system can support multiple sources of uncertainty. In contrast to all these analyses, Diamont is the only one flexible enough to simultaneously support *multiple* analyses and approximation sources, and in addition, extending these to *parallel* programs.

Static Analyses for Approximate Programs. Though multiple static analyses target approximate programs (e.g., [8, 9, 12, 23, 26, 28, 35, 37]), most relevant to Diamont is Parallely [19], which retains the limitations of the underlying static analyses requiring developers to provide bounds on loop iterations, array sizes, and number of processes. In contrast, Diamont successfully combines static and dynamic analysis and works on a real language (Go), which jointly allow for verification of much larger benchmarks. Additionally, Diamont also extends sequentialization for dynamic conditions.

Dynamic Analysis and Runtime Monitoring. DECAF [6] performs dynamic reliability verification through type inference. Our work avoids DECAF’s strict independence assumptions by adding reliabilities instead of multiplying (both bounds

Table 2: Comparison of Related Work. (\checkmark^* indicate analyses that monitor confidence intervals, which is another interpretation of Diamont’s uncertain intervals)

Method	Static	Empirical	Runtime	Reliability	Accuracy	Verified	Parallel	Multi-Source
Diamont	\checkmark	\times	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Parallely	\checkmark	\times	\times	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Rely	\checkmark	\times	\times	\checkmark	\times	\checkmark	\times	\times
Chisel	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\times	\checkmark
DECAF	\checkmark	\times	\checkmark	\checkmark	\times	\checkmark	\times	\times
EnerJ	\checkmark	\checkmark	\times	\times	\times	\checkmark	\times	\times
AffineFloat	\checkmark	\times	\checkmark	\times	\checkmark	\checkmark	\times	\times
PAssert	\times	\checkmark	\checkmark	\checkmark^*	\checkmark^*	\checkmark	\times	\times
Uncertain<T>	\times	\checkmark	\checkmark	\checkmark^*	\checkmark^*	\times	\times	\times

are close in practice). Ringenbun et al. [32] propose offline and online approaches to monitor the quality of programs, using methods such as dataflow techniques and comparison to the precise program. Diamont instead propagates uncertain intervals during both static and dynamic phases, allowing it to monitor uncertainty with greater precision. Maderbacher et al. [25] focus on precisely correcting bitflips with minimal checks. In contrast, Diamont monitors uncertainty from many sources in programs that can tolerate some error.

AffineFloat [13] and Ceres [14] provide dynamic analysis for numerical error. Herbgrind [38] locates possible sources of numerical error. These tools measure floating point roundoff errors, but have high overhead. Diamont focuses on analyzing error from casting and external sources e.g., sensors. Uncertain<T> [4] used an early form of uncertain intervals, however they use sampling to determine error. Statistical model checking tools [39] can provide statistical guarantees on program properties expressed in a temporal logic. PAssert [36] and AxProf [21] statistically verify at development time a single probabilistic assertion at the end of the program. In contrast, Diamont supports many checks at different points in the program at runtime.

8 Conclusion

The past decade brought many techniques for developing new approximations and analyzing uncertainty for specific scenarios, but much less work has been done in integrating these diverse concepts in a unifying, rigorous, and extensible framework. Diamont aims to pave the way toward that goal – it supports multiple uncertainty sources (input noise, variable-precision code, errors in communication, and unreliability in hardware), combines static analysis and dynamic monitoring, supports a significant fragment of the Go language, and operates on several emerging applications (precision agriculture, graph analytics, and media processing).

We demonstrated the benefit of our analysis and optimizations by reducing the execution overhead to 3% on average (16.3% maximum). We believe this work can serve as a starting point for sound runtime systems in domains that need to rigorously handle uncertainty, such as robotics or the Internet-of-Things.

Acknowledgements. We thank the anonymous reviewers for their useful suggestions. The research presented in this paper was supported in part by NSF Grants No. CCF-1846354, CCF-1956374, CCF-2028861, and CCF-2008883, USDA Grant No. NIFA-2024827, and a gift from Facebook.

References

1. Achour, S., Rinard, M.: Energy Efficient Approximate Computation with Topaz. OOPSLA (2015)
2. Ahmad, M., Hijaz, F., Shi, Q., Khan, O.: CRONO: A Benchmark Suite for Multithreaded Graph Algorithms Executing on Futuristic Multicores. In: IISWC (2015)
3. Bakst, A., Gleissenthall, K.v., Kici, R.G., Jhala, R.: Verifying distributed programs via canonical sequentialization. In: OOPSLA (2017)
4. Bornholt, J., Mytkowicz, T., McKinley, K.S.: Uncertain<T>: A First-order Type for Uncertain Data. In: ASPLOS (2014)
5. Boston, B., Gong, Z., Carbin, M.: Leto: verifying application-specific hardware fault tolerance with programmable execution models. In: OOPSLA (2018)
6. Boston, B., Sampson, A., Grossman, D., Ceze, L.: Probability type inference for flexible approximate programming. In: OOPSLA (2015)
7. Boyapati, R., Huang, J., Majumder, P., Yum, K.H., Kim, E.J.: APPROX-NoC: A Data Approximation Framework for Network-On-Chip Architectures. In: ISCA (2017)
8. Carbin, M., Kim, D., Misailovic, S., Rinard, M.: Proving Acceptability Properties of Relaxed Nondeterministic Approximate Programs. PLDI (2012)
9. Carbin, M., Kim, D., Misailovic, S., Rinard, M.: Verified integrity properties for safe approximate program transformations. PEPM (2013)
10. Carbin, M., Misailovic, S., Rinard, M.: Verifying Quantitative Reliability for Programs That Execute on Unreliable Hardware. In: OOPSLA (2013)
11. Chen, Y., Louri, A.: An Approximate Communication Framework for Network-on-Chips. IEEE Transactions on Parallel and Distributed Systems (2020)
12. Darulova, E., Izycheva, A., Nasir, F., Ritter, F., Becker, H., Bastian, R.: Daisy-Framework for Analysis and Optimization of Numerical Programs. In: TACAS (2018)
13. Darulova, E., Kuncak, V.: Trustworthy Numerical Computation in Scala. In: OOPSLA (2011)
14. Darulova, E., Kuncak, V.: Certifying solutions for numerical constraints. In: RV (2012)
15. de Kruijf, M., Nomura, S., Sankaralingam, K.: Relax: an architectural framework for software recovery of hardware faults. ISCA (2010)
16. Dean, J., Ghemawat, S.: MapReduce: Simplified data processing on large clusters. OSDI (2004)
17. Fernando, V., Franques, A., Abadal, S., Misailovic, S., Torrellas, J.: Replica: A Wireless Manycore for Communication-Intensive and Approximate Data. In: ASPLOS (2019)
18. Fernando, V., Joshi, K., Laurel, J., Misailovic, S.: Appendix to Diamont. <https://vimuth.github.io/diamont/appendix.pdf> (2021)
19. Fernando, V., Joshi, K., Misailovic, S.: Verifying Safety and Accuracy of Approximate Parallel Programs via Canonical Sequentialization. In: OOPSLA (2019)
20. Golubovic, N., Krintz, C., Wolski, R., Sethuramasamyraja, B., Liu, B.: A scalable system for executing and scoring K-means clustering techniques and its impact on applications in agriculture. International Journal of Big Data Intelligence **6** (2019)
21. Joshi, K., Fernando, V., Misailovic, S.: Statistical algorithmic profiling for randomized approximate programs. In: ICSE (2019)
22. Joshi, K., Fernando, V., Misailovic, S.: Aloe: Verifying Reliability of Approximate Programs in the Presence of Recovery Mechanisms. CGO (2020)
23. Lahiri, S., Haran, A., He, S., Rakamaric, Z.: Automated Differential Program Verification for Approximate Computing. Tech. rep. (May 2015)
24. Liu, T.: Datasheet for AM2302 Sensor. <https://cdn-shop.adafruit.com/datasheets/Digital+humidity+and+temperature+sensor+AM2302.pdf> (2020)

25. Maderbacher, B., Karl, A.F., Bloem, R.: Placement of Runtime Checks to Counteract Fault Injections. In: RV (2020)
26. Misailovic, S., Carbin, M., Achour, S., Qi, Z., Rinard, M.: Chisel: Reliability- and Accuracy-aware Optimization of Approximate Computational Kernels. In: OOPSLA (2014)
27. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank citation ranking: Bringing order to the web. Tech. rep. (1999)
28. Panckekha, P., Sanchez-Stern, A., Wilcox, J.R., Tatlock, Z.: Automatically Improving Accuracy for Floating Point Expressions. In: PLDI (2015)
29. Paradis, L., Han, Q.: A survey of fault management in wireless sensor networks. *Journal of Network and systems management* (2007)
30. Ranjan, A., Venkataramani, S., Fong, X., Roy, K., Raghunathan, A.: Approximate Storage for Energy Efficient Spintronic Memories. In: DAC. 2015 (2015)
31. Recht, B., Re, C., Wright, S., Niu, F.: Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In: *Advances in neural information processing systems* (2011)
32. Ringenburt, M., Sampson, A., Ackerman, I., Ceze, L., Grossman, D.: Monitoring and debugging the quality of results in approximate programs. ASPLOS (2015)
33. Rubio-González, C., Nguyen, C., Nguyen, H., Demmel, J., Kahan, W., Sen, K., Bailey, D., Iancu, C., Hough, D.: Precimonious: Tuning assistant for floating-point precision. SC (2013)
34. Samadi, M., Jamshidi, D.A., Lee, J., Mahlke, S.: Paraprox: Pattern-based Approximation for Data Parallel Applications. In: ASPLOS (2014)
35. Sampson, A., Baixo, A., Ransford, B., Moreau, T., Yip, J., Ceze, L., Oskin, M.: Accept: A programmer-guided compiler framework for practical approximate computing. Tech. rep. (2015)
36. Sampson, A., Panckekha, P., Mytkowicz, T., McKinley, K., Grossman, D., Ceze, L.: Expressing and verifying probabilistic assertions. PLDI (2014)
37. Sampson, A., Dietl, W., Fortuna, E., Gnanapragasam, D., Ceze, L., Grossman, D.: EnerJ: Approximate data types for safe and general low-power computation. In: PLDI (2011)
38. Sanchez-Stern, A., Panckekha, P., Lerner, S., Tatlock, Z.: Finding root causes of floating point error. In: PLDI (2018)
39. Sen, K., Viswanathan, M., Agha, G.: Statistical model checking of black-box probabilistic systems. In: CAV (2004)
40. Sidiroglou, S., Misailovic, S., Hoffmann, H., Rinard, M.: Managing Performance vs. Accuracy Trade-offs With Loop Perforation. FSE (2011)
41. Stanley-Marbell, P., Alaghi, A., Carbin, M., Darulova, E., Dolecek, L., Gerstlauer, A., Gillani, G., Jevdjic, D., Moreau, T., Cacciotti, M., Daglis, A., Jerger, N.E., Falsafi, B., Misailovic, S., Sampson, A., Zufferey, D.: Exploiting Errors for Efficiency: A Survey from Circuits to Applications. *ACM Computing Surveys Journal*, (2020)
42. Stevens, J.R., Ranjan, A., Raghunathan, A.: AxBA: an approximate bus architecture framework. In: ICCAD (2018)
43. TensorFlow Developers: Tensorflow (2021). <https://doi.org/10.5281/zenodo.5159865>, <https://www.tensorflow.org>
44. Yazdanbakhsh, A., Mahajan, D., Esmailzadeh, H., Lotfi-Kamran, P.: AxBench: A Multiplatform Benchmark Suite for Approximate Computing. *IEEE Design Test* **34**(2) (April 2017)
45. Zhuang, W., Chen, X., Tan, J., Song, A.: An empirical analysis for evaluating the link quality of robotic sensor networks. In: WCSP (2009)