Formal Verification of Error Correcting Circuits Using Computational Algebraic Geometry

Alexey Lvov, Luis A. Lastras-Montaño
IBM T.J.Watson Research Center
Yorktown Heights, NY, 10598
{lvov, lastrasl}@us.ibm.com

Viresh Paruthi, Robert Shadowen, Ali El-Zein
IBM Systems and Technology Group
Austin, TX, 78758
{vparuthi, shadowen, elzein}@us.ibm.com

Abstract— **Algebraic error correcting codes (ECC) are widely used to implement reliability features in modern servers and systems and pose a formidable verification challenge. We present a novel methodology and techniques for provably correct design of ECC logics. The methodology is comprised of a design specification method that directly exposes the ECC algorithm's underlying math to a verification layer, encapsulated in a tool "BLUEVERI" , which establishes the correctness of the design conclusively by using an apparatus of computational algebraic geometry (Buchberger's algorithm for Gröbner basis construction). We present results from its application to example circuits to demonstrate the effectiveness of the approach. The methodology has been successfully applied to prove correctness of large error correcting circuits on IBM's POWER systems to protect memory storage and processor to memory communication, as well as a host of smaller error correcting circuits.**

I. Introduction

ECCs are widely used in practice to protect data against random errors that inevitably occur during transmission as well as during prolonged storage. As semiconductor technology is scaling down to the nanometer regime and tens of gigabits per second transmission rates, error-free data handling requires larger and more sophisticated error correcting circuits, with the code construction and encoding/decoding algorithms almost always going beyond the templates found in classical literature due to feature set requirements. For example, the IBM z196 systems feature "RAIM" (Redundant Array of Independent Memory, [1], [2]) with a 90 byte ECC that allows the system to recover instantaneously from a full DIMM failure even in the presence of additional chip failures. Each such error correcting circuit has to be individually designed and programmed by a human designer. The resulting implementation complexity in hardware can lead to design errors which can cause costly re-spins of the Silicon and derail schedules. Establishing correctness/verification of such complex hardware is of critical importance, though poses formidable challenges.

Traditional verification methods such as software simulation, hardware-accelerated simulation or post-Silicon debug offer insufficient coverage given the difficult nature of the logic and the large solution space to be investigated. State-of-the-art formal verification algorithms (which inherently check circuit behavior against all possible legal combinations of inputs) offering high capacity have been found lacking in proving correctness because of their inability to exploit the specifics of the underlying algebra - Galois field arithmetic.

We propose a solution to the problem of complete symbolic verification of logical circuits which substantially rely on arithmetic over Galois fields. Most of the error correcting circuits fall in the above category, as well as some of the circuits for data encryption and arithmetic logic unit (ALU).

The verification technique is encapsulated in a reasoning tool "Blue Code Verifier" - "BLUEVERI" - and applies algebraic geometry methods (e.g. checks on the consistency of polynomial systems of equations using the concept of Gröbner basis and the associated Buchberger's algorithm) to the problem of verifying circuits defined over Galois fields in order to establish correctness of the logic circuit against a mathematical specification. The methodology has been successfully applied to verify real life error correcting codes at IBM resulting in substantially improved verification quality, by providing full proof of the correctness of the design which was otherwise unobtainable, and in improved productivity, via significantly reduced verification time and effort. We expect the improvements to accumulate as the methodology gets applied "out-of-the-box" to future processor chips employing even stronger ECC designs, and will be key to integrate commodity memories in products as well as in the design of communication link transceivers. The techniques involved are applicable to other types of logic circuitry based on Galois field arithmetic such as Elliptic Curve Cryptography.

A. Previous Art

Simulation-based methods such as software simulation or hardware-accelerated simulation are inapplicable to the problem of complex ECC verification. This is due to the fact that the problem has large numbers of inputs which precludes an exhaustive exploration to fully verify the ECC circuitry to cover all possible combinations of input bit strings and injected errors (within the claimed error correction capability of the code) and check to see if in each case the decoded bit string is equal to the original one. Directed simulation to cover the vast majority, if not all, of "corner cases" again requires a careful analysis of the code to enumerate correction capability and features - a process which is inherently subject to human limitations and errors. Systematic methods such as SAT or graph-based canonical representations of the logic with Decision Diagrams (DD) such as BDDs [3], BMDs [4], FDDs [5] run

978-0-9835678-2-0/12/$31.00 ©2012 FMCAD Inc.

out of steam quickly due to the large input space and the complexity of the underlying logic employing exclusive-ORs. Our experience suggests that these existing decision procedures have difficulty scaling to designs beyond circuits with more than 24-bit inputs. Enhanced verification techniques leveraging Transformation-based Verification (TBV) [11] concepts to simplify then prove the designs become capacity gated for 32-bit Galois field algorithms and beyond. Satisfiability Modulo Theory (SMT) solvers which utilize specialized theories to address specific problem domains (e.g. bit-vectors) do not address polynomial equation solving over Galois fields. Our approach addresses this niche and proposes a methodology to solve such systems of polynomial equations over Galois fields efficiently.

A search for verification of Galois field circuits reveals the following applicable references - [6] and [7]. [6] defines a formal first-order logic language for symbolic arithmetic over an arbitrary binary Galois field along with a set of rules for manipulation of formal sentences (such as transformation of the sentence into prenex normal form, usage of DeMorgan's law, elimination of variables etc.). The correctness criterion for parts of some ECC circuits can be formally expressed in this language, e.g. finding the error locator polynomial from the value of the syndrome for Reed-Solomon codes. A formal reasoning in the language is then applied to prove or disprove the correctness statement. The method is only applicable to verification of algorithms which are correct in any $GF(2^k)$ independently of the value of k. In our method the size of the field is specified; in particular this allows the use of constants of the field other than '0' and '1' in the circuit. The method does not employ any of the computational algebraic geometry machinery; that bounds it to purely $GF(2^k)$ circuits (with no bit operations allowed), while our method works on circuits with mixed bit and GF signals (Boolean result of test value operations on GF signals is computed by building Gröbner basis of polynomial algebraic system).

The latter [7] applies Gröbner basis techniques to the very narrow problem of verifying multipliers over a large Galois field. The class of the multipliers is further limited to those based on representation of the large field as an extension of degree m of a smaller field of degree n. The paper reports practical results of verifying multipliers up to maximum field size of $GF(2^{1024}), (m = 32, n = 32)$, but it does not make any attempts to verify circuits other than this multiplier circuit with a fixed structure parameterized with only two integers m and n. In contrast our method is capable of verifying virtually any circuit built with GF, Boolean and mixed operations, with the runtime and memory being the only limiting factors for large circuits.

II. Proposed Method

Our method was first inspired by the need to verify a large 1024-bit input error correction circuit responsible for protecting the memory store as well as the communication between a POWER processor and memory. A traditional formal verification approach to verify the circuitry quickly became intractable given the vast search space.

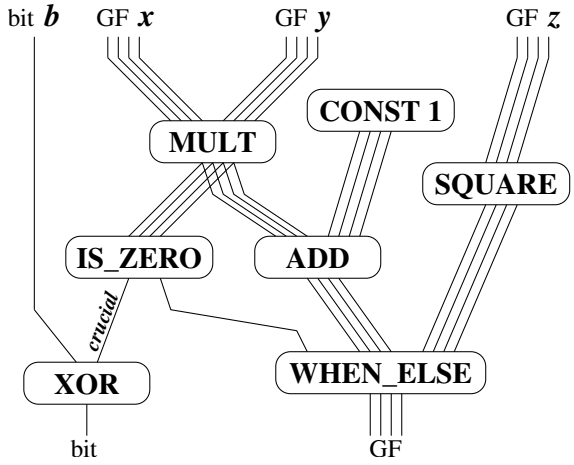

Fig. 1. Example of BLUEVERI circuit representation.

The main idea is to use the fact that algebraic ECCs operate mostly on the elements of finite fields, and there are powerful techniques for symbolic reasoning in this domain. The process of verification of such circuits reduces to the verification of a number of algebraic statements of the type "A certain system of multivariate polynomials over a finite field implies some other system of multivariate polynomials over a finite field". The latter problem relates to computational algebraic geometry and can be solved by building Gröbner bases for certain sets of polynomials by using Buchberger's algorithm ([8], pp.77, 82-87).

A. Verification Set-up

The verification set-up consists of two parts: the circuit to be verified, and a check file containing information about the set of legal inputs and the expected values for some set of "crucial" signals; an example of the latter would be an uncorrectable error flag (see subsection III-A) or a signal that tests the equality between two bit vectors (see subsection III-B). The verification task at hand is to formally prove (or disprove) that for any legal combination of inputs, the values of the crucial signals match their expected values.

In a standard processing methodology, the circuit is generally represented by a directed graph where the edges are wires carrying only Boolean signals, and nodes are gates performing only basic Boolean operations. Since we assume that a large portion of the operations in the circuit are operations in $GF(2^k)$ arithmetic, we modify this representation by "glueing" together wires which represent the same $GF(2^k)$ elements and putting "black boxes" around the pieces of the circuit which represent basic $GF(2^k)$ arithmetic operations. Practically this is done by passing a special option to the HDL compiler, telling it to not synthesize functions from a given list. The circuit in our representation typically looks similar to the example on Fig 1.

After this transformation, each wire carries either a Boolean

signal or a $GF(2^k)$ signal. For this reason, we generalize the concept of "gate" so that now each gate performs one of the following operations:

- Basic binary arithmetic operations on $GF(2^k)$: **ADD** (both $x{+}y$ and $x{-}y$), **MULT** (xy), **DIV** (xy^{2^k-2}).
- Any fixed set of unary operations on $GF(2^k)$ which are linear over $GF(2)$, e.g. Frobenius automorphism (square), projections on elements of a fixed basis, square root, bit permutations etc.
- Any fixed set of $GF(2^k)$ constants (functions without arguments).
- **WHEN_ELSE**(b, x, y) function which returns $GF(2^k)$ element x when bit b is 1 and $GF(2^k)$ element y otherwise.
- $GF(2^k)$ value test functions which return value is a bit: **IS_ZERO**(x), **IS_NONZERO**(x).
- Boolean functions: **NOT**, **AND**, **OR**, **XOR**.

The check file contains algebraic constraints on the $GF(2^k)$ inputs, optionally initial values for some Boolean and $GF(2^k)$ inputs, and the expected values for the crucial Boolean signals testing the desired behavior for the circuit. The crucial signals are restricted to Boolean because any condition on $GF(2^k)$ signals can be expressed as a condition on Boolean signals by adding just a few gates to the circuit. For example, if one wants to state that a GF signal x is equal to a given constant $const$, then one may alternatively assert that we expect

$$\big(\text{IS\_ZERO}(\text{ADD}(x, const)\big)$$

to be equal to 1.

The algebraic constraints are specified in conjunctive normal form (CNF) whose literals are multivariate polynomial equalities or inequalities on the free variables associated with each of the $GF(2^k)$ inputs.

Here is an example of a check file for the circuit on Fig 1:

```
BEGIN_CHECK;

  IN_BITS_SETTINGS;
    b <= '0';

  EXPLICIT_EXPRESSIONS_FOR_SOME_GF_INPUTS;
    x <= "8F3A";

  ALGEBRAIC_CONSTRAINTS_ON_GF_INPUTS;
    [ (y^3 + z^5 == 0) or (y^2 + z != 0) ]
    and
    [ (y == 0) or (z == 0) or (y + z != 0) ]

  BIT_EXPECTED_VALUES;
    crucial must be '1';

END_CHECK;
```

We support multiple checks in one check file in which case our tool verifies them independently one by one, and appending new checks at the end of the file during verification (a necessary feature for the "fork on unresolved bits" mechanism outlined later).

B. Verification Flow

The process starts by assigning a free variable (e.g. the symbolic string identifier used in the HDL file) to each of the $GF(2^k)$ inputs. Next the values of the crucial bit signals are computed one by one by applying the following recursive procedure. The procedures for "...execute the operation ..." will be explained for each type of operation subsequently.

```
COMPUTE_OUTPUT_OF_GATE(signal g) {
  // case g is Boolean : Attempt to compute to const. '0' or '1' .
  // case g is GF(2^k) : Compute as a symbolic rational
                         expression in the free variables.
  for all inputs g_i of g {
    COMPUTE_OUTPUT_OF_GATE(g_i)
  }
  switch (type of g) {
    ADD:  ...Execute the operation ...
    MULT: ...Execute the operation ...
    ...  ...
    XOR:  ...Execute the operation ...
  }
}
```

Given unlimited time and memory and assuming that all recursive sub-calls successfully compute values of $g_1, g_2, \ldots$ a call to COMPUTE_OUTPUT_OF_GATE(g) always succeeds if g is a $GF(2^k)$ signal. However, it may fail for Boolean signals because Boolean signals are (generally) not constants but depend on the inputs. If a Boolean signal cannot be computed to '0' or '1' we skip to the next check and add two new checks at the end of the check file assuming values '0' and '1' for that bit by applying the "fork on unresolved bit" procedure described later in this subsection. Note that although it may seem that this would fork on nearly every bit in the circuit, in our experience for ECCs the situation is typically just the opposite: given a restricted set of inputs (e.g. exactly one injected error) most of the Boolean signals in the circuit do not depend on the inputs; an example of this can be seen in subsection III-A in the computation of the uncorrectable error flag of a decoder [1]. Furthermore, BLUEVERI performs signal dependency checks that result in the value of many boolean signals in the circuit not being needed; such booleans never cause a fork as described above.

Given $g_1, g_2, \ldots$, we compute g depending on the type of operation as follows:

ADD and **MULT** : Perform the operation on the mutivariate rational expressions. E.g. ADD($\frac{x}{y+z}$, $\frac{y}{x+z}$) = $\frac{x^2+xz+y^2+yz}{xy+xz+yz+z^2}$, MULT($x+1$, $y+1$) = $xy+x+y+1$ etc.

UNARY_LINEAR_i : Any operation on $GF(2^k)$ which is linear over $GF(2)$ can be given by a linearized polynomial (a polynomial containing only terms of the form cx^{2^t}, see [9] pp.107-124). Substitute the input rational expression into the linearized polynomial. E.g. in $GF(16)$ $\text{Tr}(x) \stackrel{\text{def}}{=} x^8+x^4+x^2+$

[1]Very often the uncorrectable error signal is both an internal signal upon which further things depend and also an output by itself.

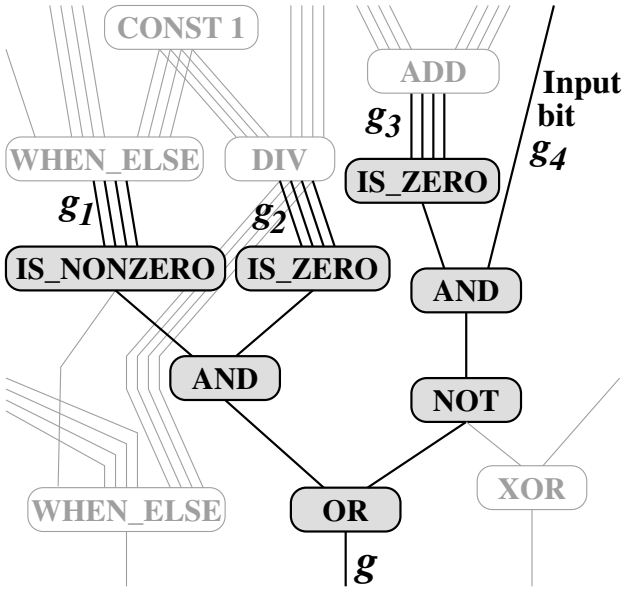

Fig. 2. Example of maximal "algebraic system" subgraph for signal g.

x, $\mathrm{Tr}(y+z^3) = y^8+y^4+y^2+y+z^{24}+z^{12}+z^6+z^3$.

CONST_i: Set signal g to the constant (a rational expression containing no free variables).

WHEN_ELSE(b, X, Y): Set rational expression g to rational expression X if b is '1' and to rational expression Y otherwise.

IS_ZERO, **IS_NONZERO**, **NOT**, **AND**, **OR**, **XOR**: Computation of values of gates with Boolean output constitutes the most complex part of our algorithm.

To compute the value of g we first find the maximal subgraph consisting of all gates h_j such that there exists a directed path from h_j to g and all gates on this path except for h_j itself are elementary Boolean gates (NOT, AND, OR or XOR). An example is shown on Fig. 2. Note that the subgraph may only contain IS_ZERO, IS_NONZERO and elementary Boolean gates, and any IS_ZERO or IS_NONZERO in the subgraph must be a top most gate. The input signals g_i of the subgraph are either $GF(2^k)$ inputs of value test functions or Boolean inputs of the whole circuit.

By inductive hypothesis for our recursive function COMPUTE_OUTPUT_OF_GATE(g) all $GF(2^k)$-type g_i have already been assigned some rational expression in the free variables, and all Boolean type g_i have been computed to constant '0' or '1' (this is possible for all Boolean inputs to the circuit due to an explicit assignment in the "In bits settings" section of the check which may be set either by the user or as a result of forking on unresolved bits).

The Boolean function given by the subgraph can be written as a conjunctive normal form whose literals are $g_i = 0$ or $g_i \neq 0$, where g_i are rational expressions. As we will show in the description of DIV operation, we always make sure the denominators of our rational expressions cannot be zero. This allows replacement of $g_i = 0$ and $g_i \neq 0$ literals by $\mathrm{numerator}(g_i) = 0$ and $\mathrm{numerator}(g_i) \neq 0$ polynomial equalities/inequalities and express g as an algebraic system of the form

$$\begin{cases} \big[P_*(x_0,x_1,\ldots) =, \neq 0\big] \text{or} \ldots \text{or} \big[P_*(x_0,x_1,\ldots) =, \neq 0\big], \\ \quad\ldots \quad\quad \ldots \\ \big[P_*(x_0,x_1,\ldots) =, \neq 0\big] \text{or} \ldots \text{or} \big[P_*(x_0,x_1,\ldots) =, \neq 0\big], \end{cases} \quad (1)$$

where P_* denote arbitrary polynomials in the free variables $x_0, x_1, x_2, \ldots$ associated with the $GF(2^k)$ inputs of the circuit.

The algebraic constraints on the inputs are also given as CNF, and form an algebraic system of the same type.

g is constant'0' if and only if

{input constraints CNF} AND {g-subgraph CNF} (2)

is unsatisfiable.

g is constant'1' if and only if

{input constraints CNF} AND NOT{g-subgraph CNF} (3)

is unsatisfiable.

Each of the expressions (2) and (3) can be converted to a single CNF of the form (1). Hence, it suffices to show how to check whether a system of the form (1) is unsatisfiable.

Satisfiability checking algorithm:

The first step is to get rid of inequalities in the system. For each inequality $P_*(x_0, x_1, \ldots) \neq 0$ we introduce an auxiliary free variable t_* and replace the inequality by

$$t_* \cdot P_*(x_0, x_1, \ldots) - 1 = 0.$$

One can easily check that if the system before replacement is satisfiable in variables $\{x_0, x_1, \ldots, t_0, t_1, \ldots\}$ then the system after replacement is satisfiable in variables $\{x_0, x_1, \ldots, t_0, t_1, \ldots\} \cup \{t_*\}$ and vice versa.

The new system contains only polynomial equalities. Next we replace all OR operations with multiplication:

$$\begin{cases} \big(Q_*(\{x_*\},\{t_*\})\big) \cdot \ldots \cdot \big(Q_*(\{x_*\},\{t_*\})\big) = 0, \\ \quad\ldots \quad\quad \ldots \\ \big(Q_*(\{x_*\},\{t_*\})\big) \cdot \ldots \cdot \big(Q_*(\{x_*\},\{t_*\})\big) = 0, \end{cases}$$

Now we have a regular algebraic system of multivariate polynomials over $GF(2^k)$.

By Hilbert's Weak Nullstellensatz a system of multivariate polynomials is unsatisfiable over an algebraically closed field if and only if the ideal generated by the polynomials of the system coincides with the whole ring (i.e. contains 1) (refer [8], pp. 169-173), $x \in GF(2^k)$ if and only if $\big[x \in \mathrm{alg\_closure}\big(GF(2^k)\big)$ AND $x^{2^k} - x = 0\big]$. For each variable v_* of our system add equation $v_*^{2^k} - v_* = 0$. The new system (denote it S) is satisfiable in the algebraic closure of $GF(2^k)$ if and only if the original system is satisfiable in $GF(2^k)$.

Next we build a Gröbner basis of the ideal given by the polynomials of system S. This can be done by Buchberger's algorithm ([8], pp. 77, 82-87). The original system is unsatisfiable in $GF(2^k)$ if and only if the Gröbner basis of S contains 1.

If the value of g is proved to be a constant '0' or '1' assign this value to g (computation successful). Otherwise fork on the unresolved Boolean signal g as follows:

Add two copies of the current check at the end of the check file as given below.

- If g is an input Boolean signal add `g <= '0'` to the "In bits settings" section of copy_1 and `g <= '1'` to the "In bits settings" section of copy_2.
- Otherwise add NOT(System (1)) to the conjunctive normal form in "Algebraic constraints on GF inputs" section of copy_1 and System (1) to the CNF in "Algebraic constraints on GF inputs" section of copy_2.

Skip the current check and continue to the next one with the two additional checks added at the end of the queue. As a side note, the two examples in subsections III-A and III-B do not require branching of this type for completion.

The only operation we have not explained yet is division. **DIV** : In logical circuits division is usually implemented as `if` $y \neq 0$ `return x/y; else return 0;` (which is equivalent to xy^{2^k-2}). To compute the result of division we first attempt to prove that the constraints on the inputs imply that the divisor is either always $= 0$ or always $\neq 0$ by the same algebraic method as for the gates with Boolean output. If successful, we simply assign 0 or the rational $expressionx/y$ to g. Otherwise we fork on the test of [denominator $= 0$] the same way as shown above for non-input Boolean signals.

We have shown how to compute value of any gate given the values of its inputs. $GF(2^k)$ signals are computed as symbolic rational expressions in the input signals, and Boolean signals must compute to constant '0' or '1' creating new branches with additional algebraic constraints on the inputs if necessary. This completes the description of our algorithm.

Our actual C implementation contains many more features than described above. The most important ones include:

- Careful manipulations of conjunctive normal form systems: A brute force manipulation of CNFs, and opening parenthesis in polynomial products which come from large OR-clauses would cause an immediate exponential explosion of the size of the system. However special care is taken of systems of the form (1) which most commonly appear in algebraic circuits. This prevents a rapid increase of the size of the system - at least for typical cases. In particular, if g-CNF has only one OR clause of length ≥ 2, i.e. has the form $\Big([P_* =, \neq 0] \text{ or } \ldots \text{ or } [P_* =, \neq 0]\Big)$ and $[P_* =, \neq 0]$ and $\ldots$ and $[P_* =, \neq 0]$, our implementation ensures the size of any system for which we build a Gröbner basis is simply equal to the sum of the sizes of the input constraints system and g-CNF system.
- "Lazy" signal computation method: In order to find values of expressions such as ('1' or x), ('0' and x), (when '1' : $const$ else x) etc., we do not compute x. This gives a significant speed up especially when the signals whose values we need to verify are localized in a relatively small part of a large circuit.
- Verification flow control: The user can control a number of verification process options such as whether to spend more time on Gröbner basis computation of a given bit vs. fork; whether to attempt to save time by skipping the $x \in GF(2^k)$ constraints which makes false negatives (but not false positives) possible; etc.

The verification process can have three possible outcomes:

1) For all checks all crucial bit values are computed and match the expected values.
2) One of the checks (including checks added by "fork on unresolved bit") fails because the value of one of the crucial bits is opposite to the expected value specified in the check file.
3) One of the checks (including checks added by "fork on unresolved bit") fails to compute one of the crucial bit values due to insufficient time or memory.

In the latter two cases an interactive bug tracing interface allows the user to browse the graph of signals and view their values in the form of symbolic rational expressions and algebraic systems.

III. Experimental Results

If there is no restriction on time and memory the verification process is guaranteed to prove or disprove the specification in the check file. We will give in what follows two simple examples (subsections III-A and III-B) where this is accomplished within a reasonable amount of time, demonstrating the power of reasoning at the Galois field level as opposed to the Boolean level. For complex, real-life designs (as exemplified in subsection III-C) we have found it useful to help BLUEVERI by manually partitioning the search space, resulting in very little use of the "forking" feature described earlier. In addition, in some instances care is taken to specify the circuit in otherwise equivalent forms to aid BLUEVERI in keeping down the size of its internal rational expressions and the complexity of algebraic systems it generates; this was not necessary in the two examples below.

A. The uncorrectable error flag of a sample Reed-Solomon decoder

As a first example, we consider a Reed-Solomon code with symbols belonging to a finite field $GF(q)$ with $q = 2^k$ elements for some integer k. We shall assume that the length of this code is $n = 2^k - 1$. Let r denote the number of check symbols of the Reed-Solomon code. We assume that this Reed-Solomon code has been furnished with a decoder that is capable of correcting any one symbol error, and can detect up to $r-1$ different errors. This decoder has a number of different components, one of which is responsible for the computation of the uncorrectable error flag. This flag is a single Boolean output that is raised whenever the decoder has detected 2,3, or up to $r-1$ errors, and kept low whenever the error scenario corresponds to a single error, or alternately whenever there is no error.

For our choice of Reed-Solomon code, the r syndromes of this Reed Solomon code can be computed from a (potentially

corrupted) encoded vector $v \in GF(q)^n$ using the formula

$$S_i = \sum_{j=0}^{n-1} v_j \omega^{ij}$$

for $i \in \{0, \cdots, r-1\}$, where ω denotes a primitive element of the field. Furthermore, letting $e \in F_q^n$ denote the *error vector* affecting v, so that $v = e + x$ where $x \in F_q^n$ is the uncorrupted codeword it is also known that due to linearity of the addition operator in finite fields and the vector that x has zero syndrome,

$$S_i = \sum_{j=0}^{n-1} e_j \omega^{ij} \tag{4}$$

The design of the uncorrectable error flag for this scenario is a well understood problem; for the sake of demonstration we deduce what might be a reasonable method to test it directly through formal methods. It can be easily seen from (4) that if there is only one error in e then the syndromes satisfy the following condition: $S_i S_{i+2} = S_{i+1}^2$ for $i = 0, \cdots, r-3$. Furthermore it is also known whenever e has at least one error and at most r errors, one or more of the $\{S_i\}_{i=0}^{r-1}$ is nonzero. This leads to the conjecture that one can compute the uncorrectable error flag through the following code, written using BLUEVERI VHDL style semantics:

```
t_comp : for i in 0 to r-3 generate
  t(i) <= add(mult(s(i),s(i+2)),square(s(i)));
end t_comp;
snz <= is_nz(s(0)) or ... or is_nz(s(r-1));
tnz <= is_nz(t(0)) or ... or is_nz(t(r-3));
UE  <= snz and tnz;
```

As written above, `snz` and `tnz` represent two distinct systems of equations which BLUEVERI will treat independently of each other. On the other hand, BLUEVERI will attempt to establish whether `tnz` (for example) is true or false by examining the properties of `t(0) ... t(r-3)` simultaneously as opposed to testing whether each `t(i)` is zero or not individually.

In order to test the ability of a model checker to prove the correctness of this implementation of the uncorrectable error flag, we assume that the syndrome generation portion of the decoder has been proved correct separately; this task is in fact generally computationally simpler than the one currently at hand. We then build a module that accepts inputs `e_m(0)...e_m(t-1)` (for the error magnitudes) and inputs `l(0)...l(t-1)` (for the error locations) where t is the maximum number of errors one can inject into the decoder during the test; in this particular example for the uncorrectable error flag to be correct it is known that $t = r - 1$. This module emulates the syndrome generator and computes `s(0)...s(r-1)` using the equation `s(i)` $= \sum_{i=0}^{t-1}$ `l(i) e_m(i)` (as per Equation 4), and then passes the resulting syndromes to a module that computes the uncorrectable error flag as previously described.

In order to test a variety of error scenarios, we can place constraints on `e_m(i)` and `l(i)`. For example, one can restrict the test to have exactly two errors by specifying the following constraints:

```
e(0) != 0, e(1) != 0, l(0) != 0, l(1) != 0
add(l(0),l(1)) != 0, e(1) = ... = e(t-1) = 0
```

Note that in a field of characteristic 2, addition is equivalent to subtraction, and hence the addition constraint effectively constrains `l(0) != l(1)`. These constraints can be specified in a BLUEVERI check file as equal/not equal to zero conditions on multivariate polynomial expressions. When BLUEVERI examines the dependencies of the `UE` signal, it finds that it depends on `snz` and `tnz`. BLUEVERI must either resolve that both are true, or that at least one of them is false. As described earlier, this is accomplished by invoking an attempt to compute the Gröbner basis of various system of equations related to the constraints and the expressions defining `snz` and `tnz`. Similar experiments can be conducted by updating the constraints to specify "at least two, but not more than y errors" where y is a number between 2 and $r-1$.

In order to test the capability of BLUEVERI as applied to this problem and contrast it with that of a formal prover (we chose SixthSense, IBM's state-of-the-art formal and semi-formal verification tool set, for that purpose), we set up a test with $r = 8$, $b = 8$ and with the capability to inject from 2 up to 7 errors at arbitrary locations, since the corresponding Reed-Solomon decoder is supposed to be able to detect all those errors. We also set up a parallel test with $b = 4$ which is a considerably simpler problem for a Boolean oriented formal verification system such as SixthSense [11]. The SixthSense and BLUEVERI experiments do not have any special tuning of the VHDL or the tool to improve the outcomes.

We refer the reader to Table I where the experiments were performed in a single processor (POWER6 processor @ 5GHz running AIX) and the SixthSense was run as a single software thread mainly orchestrating redundancy removal and SAT algorithms. In this set of experiments, BLUEVERI was configured to reason about the circuit with the variables (due to inputs or constraints) belonging to the *algebraic closure* of the fields. This in essence means that we did not constrain the variables to belong to the field $GF(256)$ (resp. $GF(16)$) depending on whether the symbols used were 8 bit (resp. 4 bit) symbols. The consequence of this is that although the BLUEVERI results are listed under 8-bit column, they in fact hold *for any field size*, including larger field sizes which would be even harder for a bit-level verification system to handle. Both formal systems were able to prove the correctness of the uncorrectable error flag under the single error scenario quite easily, but SixthSense was not able to prove the correctness of this flag in the double error case in the amount of time indicated in the table. In order to test the sensitivity of SXS to the field size, we performed a similar experiment for a Reed-Solomon code defined over $GF(16)$. In this case we saw better results from SixthSense, since we were able to prove the correctness of double and triple error detect cases but not four error case. It is worth noting that the field size determines many important properties of an error control code, including

symbol errors	expected UE	8 bit symbols			4 bit symbols	
		BLUEVERI	input bits	SXS	input bits	SXS
1	false	Success after 0.1 s.	16	Success after 14 s.	8	Success after 0.7 s
2	true	Success after 1 s.	32	Gives up after 24 h.	16	Success after 3 s
3	true	Success after 1 s.	48	N/A	24	Success after 55 m
4	true	Success after 33 m.	64	N/A	32	Gives up after 24h
5	true	Gives up after 6 h.	80	N/A	40	N/A

TABLE I

EXPERIMENTAL RESULTS FOR THE FORMAL VERIFICATION OF THE UNCORRECTABLE ERROR FLAG OF A SINGLE ERROR CORRECT, MULTIPLE ERROR DETECT REED-SOLOMON DECODER. SXS REFERS TO SIXTH SENSE, A BIT-LEVEL FORMAL VERIFICATION TOOL SET DEVELOPED AT IBM.

errors	8 bit symbols			4 bit symbols	
	BLUEVERI	input bits	SXS	input bits	SXS
2	Succ. 2 s.	32	Gives up after 24h	16	Succ. 0.6s
3	Succ. 2.1 s.	48	N/A	24	Succ. 16m
4	Succ. 2.1 s.	64	N/A	32	Gives up after 24h
5	Succ. 2.3 s.	80	N/A	40	N/A
6	Succ. 3.1 s.	96	N/A	48	N/A
7	Succ. 49.4 s.	112	N/A	56	N/A
8	Succ. 8m	128	N/A	64	N/A
9	Succ. 53m	144	N/A	72	N/A

TABLE II

EXPERIMENTAL RESULTS FOR THE FORMAL VERIFICATION OF THE ERROR MAGNITUDE COMPUTATION STAGE OF A REED-SOLOMON CODE.

the total codeword length, and thus it cannot be modified for the purposes of formal verification since the resulting code is entirely different and, in all likelihood, not applicable to the original problem.

B. Computing error magnitudes in a Reed-Solomon code

One of the tasks that an error control decoder for a code defined over multibit ($q > 2$) symbols must perform is to compute the locations of the symbols in error and then to compute the multibit pattern that one must add to those locations in order to correct the codeword. This multibit pattern is called the *error magnitude*. Suppose that there are t errors in a codeword, and let $\mathtt{s(0)}, \cdots, \mathtt{s(t-1)}$ be the first t syndromes (note that this example is for a different setting than the example in the previous subsection). From (4), we can derive that error magnitude computation can be carried over using the equation

$$\begin{bmatrix} \mathtt{e\_m(0)} \\ \vdots \\ \mathtt{e\_m(t-1)} \end{bmatrix} = \begin{bmatrix} 1 & \cdots & 1 \\ \mathtt{l(0)} & \cdots & \mathtt{l(t-1)} \\ \vdots & \ddots & \vdots \\ \mathtt{l(0)}^{t-1} & \cdots & \mathtt{l(t-1)}^{t-1} \end{bmatrix}^{-1} \begin{bmatrix} \mathtt{s(0)} \\ \vdots \\ \mathtt{s(t-1)} \end{bmatrix}$$

The inverse matrix above can be derived analytically. It is well known that the inverse is non singular if and only if the locations $\mathtt{l(i)}$ are all distinct of each other. This restriction can be specified through $\binom{t}{2}$ constraints each of which is a polynomial with two monomials. We refer the reader to Table III-B where we show that in this case, BLUEVERI was able to show the correctness of the corresponding circuit with up to 8 errors, while SixthSense was unable to finish the double error case within the time allocated. As in the previous subsection, in this particular example the result for BLUEVERI is actually field size independent since it exploits only the algebraic properties of the symbols. It is worth noting that the Gröbner basis machinery in BLUEVERI does get involved in proving the correctness of this circuit. This is because the inversion of the Vandermonde matrix results in rational expressions (as opposed to plain polynomial expressions) whose denominator could be zero. The task of Gröbner in here then is to show that the denominator is not zero given the assumptions on the inputs, so that BLUEVERI can proceed with the corresponding algebraic simplifications leading to the desired result.

C. A note on a real life application of BLUEVERI

The examples in the previous subsections are meant to illustrate the capabilities of a formal verification system such as BLUEVERI when compared to Boolean oriented systems. In our experience, the implementation of a real-life encoder/decoder employs many custom algorithm variants as one tries to address problems that are specific to the application at hand. In the most significant application of BLUEVERI so far, we have succeeded in proving the correctness of an ECC of a POWER microprocessor that is based on the mathematics of Reed-Solomon codes. The correctness criteria included all correctable and uncorrectable cases for which we had given guaranteed behavior (e.g. recovery from complete chip failures and detection of multiple errors). The ECC, from the decoders perspective, had over 1000 bits of input including several tens of bits worth of configuration parameters. The number of syndrome bits produced by the decoder was over 100 bits, although our testing did include testing the behavior of the encoder with analytically generated symbolic syndromes, it was not limited to it - approximately half of the total testing time exercised the more than 1000 bits of input of the circuit directly. The number of Galois field and Boolean elements in the corresponding graph is over 100,000 (compared to at most a few hundred in the previous experiments). Because of the complexity of the problem, we had to case-split to create 1M different tests, each of which exercised formally a particular region of the test space. It took about 2 weeks to prove the correctness of the entire design in a 10 machine Linux (x86) cluster.

IV. TECHNICAL SOLUTIONS

The BLUEVERI tool leverages IBM's existing front-end and simulation tools and flows. For language processing we are using Portals, IBM's HDL compiler, which accepts the

synthesizable subset of standard VHDL and Verilog languages. Portals performs behavioral synthesis on procedural HDL and produces an elaborated netlist, for BLUEVERI this is in the DADB logic database. DADB is a box-pin-net logic database used for verification flows, such as topology checking and simulator model build, which supports client transforms via a dynamically loaded plugin architecture.

Portals was modified for BLUEVERI to support the black-boxing of function calls, enabling the logic to be represented in a form amenable to analysis by BLUEVERI. High level language constructs which are output by Portals into the netlist, such as case statements, can be synthesized into lower level representations by the use of DADB client transforms.

The BLUEVERI analysis tool has its own custom input netlist format. A netlist translator was built as a DADB client to enable the tools flow from Portals into BLUEVERI.

The MESA logic simulator is a high performance cycle simulator used for functional verification within IBM. MESA simulation models are built from logic netlists in DADB by using model build clients.

The BLUEVERI code is written in C. For the computation of Gröbner bases we use "SINGULAR" [10] a powerful program for algebraic geometry computations distributed under general public license. BLUEVERI runs SINGULAR as a child process and uses "EXPECT.h", (a standard C library), for sending queries and receiving results from SINGULAR's Gröbner basis engine. The Gröbner basis obtained is for the inverse degree lexicographical ordering.

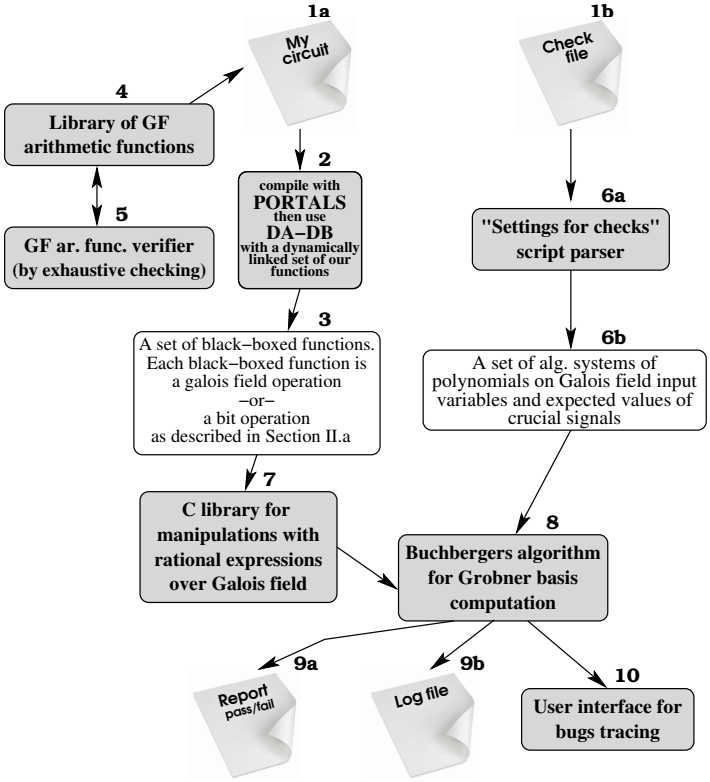

Fig. 3. General schema of BLUEVERI tool.

V. Conclusions

In this article we presented a novel technique for designing and verifying circuits based on the mathematics of Galois fields. At the heart of our approach is the idea of exposing operations on Galois field directly to a verification layer (encapsulated in a tool called BLUEVERI) which leverages powerful techniques from algebraic geometry to reason about the properties of the abstract Galois field rational expressions generated in the circuit. Our circuits are specified using a subset of existing Hardware Description Languages and as such, remain fully synthesizable, an important attribute to reduce the possibility of human error in the design process.

We demonstrated the value of the ideas we proposed in the context of two problems representative of the type of situations encountered when designing error correcting codes. In both instances, we showed BLUEVERI can significantly outperform conventional bit-level formal verification. We outlined a successful application of the BLUEVERI system to prove correctness of a real production complex error correcting code implemented on a POWER microprocessor which otherwise could not be verified conclusively with traditional verification methods.

VI. Acknowledgments

The authors would like to thank Barry M Trager, Shmuel Winograd and Geert Janseen of IBM Research for insightful discussions to help shape the solution.

References

[1] Meaney, P. J. and Lastras-Montaño, L. A. and Papazova, V. K. and Stephens, E. and Johnson, J. S. and Alves, L. C. and O'Connor, J. A. and Clarke, W. J., *IBM zEnterprise redundant array of independent memory subsystem* IBM Journal of Research and Development, Jan-Feb, Vol. 56, 2012.

[2] Lastras-Montaño, L.A.; Meaney, P.J.; Stephens, E.; Trager, B.M.; O'Connor, J.; Alves, L.C., *A new class of array codes for memory storage*, Information Theory and Applications Workshop (ITA), 2011 , vol., no., pp.1-10, 6-11 Feb. 2011

[3] R. E. Bryant *Graph Based Algorithms for Boolean Function Manipulation*, IEEE Trans. on Computers, vol. C-35, pp. 677691, August 1986.

[4] R. E. Bryant and Y-A. Chen *Verification of Arithmetic Functions with Binary Moment Diagrams*, Design Automation Conference 1995.

[5] U. Kebschull and W. Rosentiel *Efficient graph-based computation and manipulation of functional decision diagrams*, European Conference on Design Automation, pp. 278 - 282, 1993.

[6] S. Morioka, Y. Katayama and T. Yamane *Towards Efficient Verification of Arithmetic Algorithms over Galois Fields*. Proc. Computer Aided Verification 2001, vol. 2102, pp.465-477.

[7] Jinpeng Lv, Priyank Kalla and Florian Enescu *Verification of Composite Galois Field Multipliers over $GF(2^{m^n})$ Using Computer Algebra Techniques*. Proc. IEEE International High Level Design Validation and Test Workshop 2011, pp.136-143.

[8] David Cox, John Little, and Donald O'Shea. *Ideals, Varieties and Algorithms*. Undergraduate Texts in Mathematics. Springer, 2010. ISBN: 0-387-35650-9.

[9] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Encyclopedia of Mathematics and Its Applications, Volume 20. Cambridge University Press, 1997. ISBN: 0-521-39231-4.

[10] Decker, W.; Greuel, G.-M.; Pfister, G.; Schönemann, H.: SINGULAR 3-1-3 — A computer algebra system for polynomial computations. http://www.singular.uni-kl.de (2011).

[11] H. Mony, J. Baumgartner, V. Paruthi, R. Kanzelman, and A. Kuehlmann, *Scalable automated verification via expert-system guided transformations*, Formal Methods in Computer-Aided Design, 2004, pp. 159173.