

Evolutionary Fault-Tolerant Systems

Garrison W. Greenwood, Ph.D., P.E.
Dept. of Electrical & Computer Engineering
Portland State University
Portland, OR, USA 97207

All systems eventually fail. If the system is critical, it must be either repaired or replaced.

Unfortunately, repair and/or replacement is not always easy...

e.g., consider the mars rovers

In such cases, one solution is to make the system *fault-tolerant*—i.e., able to autonomously repair itself to restore functionality.

These fault-tolerant systems must (a) detect the failure, and (b) do something to fix the problem

We are going to use *evolvable hardware* (EHW) to do the fixing.

OUTLINE

I. Introduction to fault-tolerant systems

II. What is “evolvable hardware”

III. Using EHW to recover from faults

IV. Overview of real-time systems

V. Recovery under real-time constraints

So, what does fault-tolerance deal with?

Suppose a system suddenly fails. Two questions:

1. can you find the problem?

(fault detection)

2. can you fix the problem? (fault recovery)

Oh, and by the way, do both detection and recovery without human intervention...

Some observations

- fault detection is not always easy
- redundancy is the most common recovery method
- when redundancy is impractical, reconfiguration is worth considering
- in most cases detection/recovery operations have time constraints

Before discussing recovery, we need to talk about detection.

But first, “faults” and “failures” are not the same thing.

def. (failure)

inability to accomplish an assigned task

def. (fault)

a defect that leads to a failure

e.g., a clamping diode shorts (fault) which causes an input line to be permanently grounded (failure)

def. (failure modes)

A specified way in which a system or a component can fail.

e.g., the failure modes of a diode are “open” and “short”

How can we find these failure modes?

1. historical data
2. from tests or experiments
3. technical literature (journals, reports, etc.)

Once the component failure modes are known, the system are found using

Fault Tree Analysis (FTA)

A top-down approach where a system failure is assumed to have happened and one tries to find the fault that caused it.

Failure Modes & Effects Analysis (FMEA)

A bottom-up approach where the effect of every failure mode of every component is determined and analyzed.

So what causes faults?

1. components have limited lifespans (e.g., connectors corrode and transistors burn out.)
2. operational environment changes

An observation...

Environmental conditions include humidity, temperature, shock vibration and radiation.

Unanticipated environmental changes are of most concern because it makes the system operate in an unpredictable manner.

def. (fault-tolerant system (FTS))

A system that continues to operate in the presence of failures)albeit with degraded performance)

FT = fault detection + fault recovery

FT can be achieved by

- fault masking (doesn't fix the problem, just hides it)
- fault recovery via redundancy (most common method)
- fault recovery via reconfiguration (the EHW approach)

IMPORTANT NOTES:

1. In some cases only degraded performance may be achieved after fault recovery takes place
2. When a system fails, and the cause is not for environmental reasons, redundancy is the only recovery method guaranteed to restore full functionality.
3. If redundancy is not possible, the only viable recovery method may be reconfiguration.

EVOLVABLE HARDWARE

$$\text{EHW} = \text{EA} + \text{reconfigurable circuitry}$$

FPGA, FPAA, and FPTA are reconfigurable devices

“EA” means ES, EP, GA, etc.

$$P(t + 1) = \mathcal{S}(\mathcal{E}(\mathcal{V}(P(t))))$$

$P(t)$ population of solutions at time t

$\mathcal{V}(\cdot)$ random variation operator

$\mathcal{E}(\cdot)$ evaluation operator

$\mathcal{S}(\cdot)$ selection operator

Each “solution” is a hardware configuration

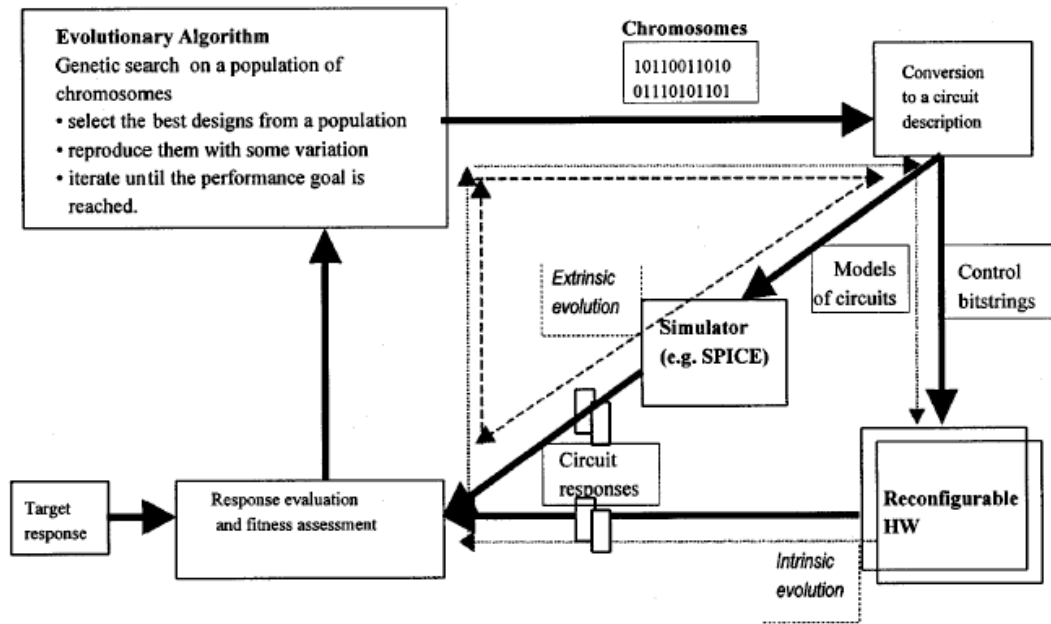
Evaluation can be done two ways

1. in hardware (*intrinsic*)

every configuration is physically implemented in hardware and then tested to determine its performance

2. in software (*extrinsic*)

every configuration is evaluated using a simulator. Only the final best solution is ever physically implemented in hardware.



Main steps for the evolutionary synthesis of electronic circuits.

From A. Stoica et al., "Reconfigurable VLSI Architectures for Evolvable Hardware: From Experimental Field Programmable Transistor Arrays to Evolution-Oriented Chips", *IEEE Trans. on VLSI Sys.*, Vol 9(1), 2001

Field Programmable Gate Array (FPGA)




Figure 1: Basic FPGA Block Diagram

from Xilinx Spartan family datasheet, Version 1.4, Jan 1999

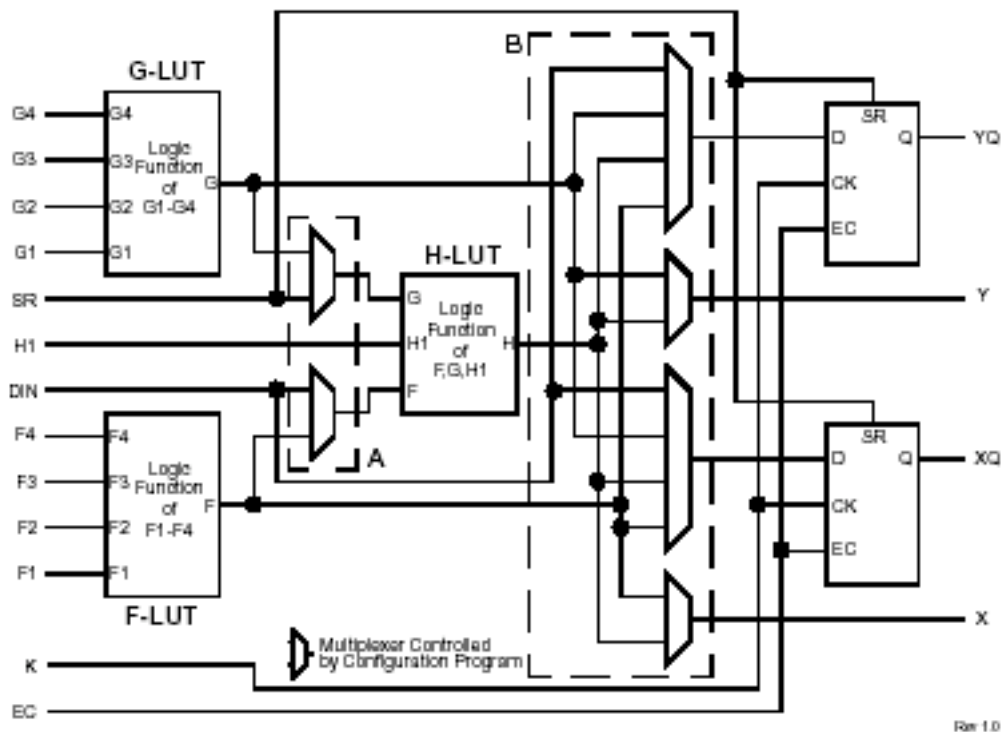


Figure 2: Spartan Simplified CLB Logic Diagram (some features not shown)

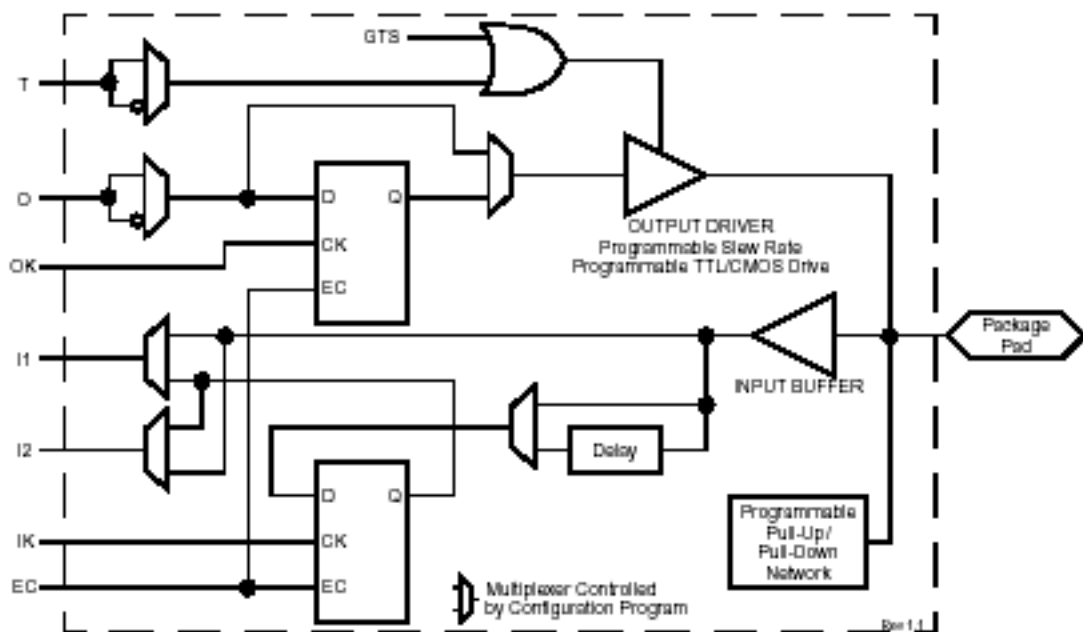


Figure 5: Simplified Spartan IOB Block Diagram

figures from Xilinx Spartan family datasheet, Version 1.4, Jan 1999




Figure 8: Spartan Series CLB Routing Channels and Interface Block Diagram




Figure 10: Programmable Switch Matrix

figures from Xilinx Spartan family datasheet, Version 1.4, Jan 1999

So, how do you program these FPGAs with an EA?

A good example is the use of “Jbits” used to program the Xilinx Virtex FPGA family.

Jbits are a set of Java classes that provide an interface into the Virtex FPGA configuration bitstream. (The bitstream can come from a design tool or readback from the FPGA itself.)

Each configurable logic block (CLB) in the FPGA has coordinates (i, j) and each CLB configuration is located in a specific region of the datastream.

The EA creates a new configuration and then used the Jbit interface to modify the bit-stream accordingly. The modified bitstream is reloaded back into the FPGA.

G3	G2	G1	G0	OR	AND
0	0	0	0	0	0
0	0	0	1	1	0
0	0	1	0	1	0
0	0	1	1	1	0
0	1	0	0	1	0
0	1	0	1	1	0
0	1	1	0	1	0
0	1	1	1	1	0
1	0	0	0	1	0
1	0	0	1	1	0
1	0	1	0	1	0
1	0	1	1	1	0
1	1	0	0	1	0
1	1	0	1	1	0
1	1	1	0	1	0
1	1	1	1	1	1

To change a LUT implementation of an OR gate to an AND gate, the bitstream is uploaded the bits of the LUT are located, and

OR 1111 1111 1111 1110 0xfffe

is changed to

AND 1000 0000 0000 0000 0x8000

```

/* Attempt to connect to the hardware */
result = board.connect(remoteHostName,
                       port);

/* Get the type of devices used by the
** first FPGA on the board */
deviceType = board.getDeviceType(0);

/* Read in the bit file */
jBits.read(infileName);

/* Set the Top left CLB (1,1) to an AND gate
** i.e. 1000 0000 0000 0000 */
jBits.set(1, 1, LUT.SLICE0_G, 0x8000);

/* Get the now updated Bitstream */
bs = jBits.getAllPackets();

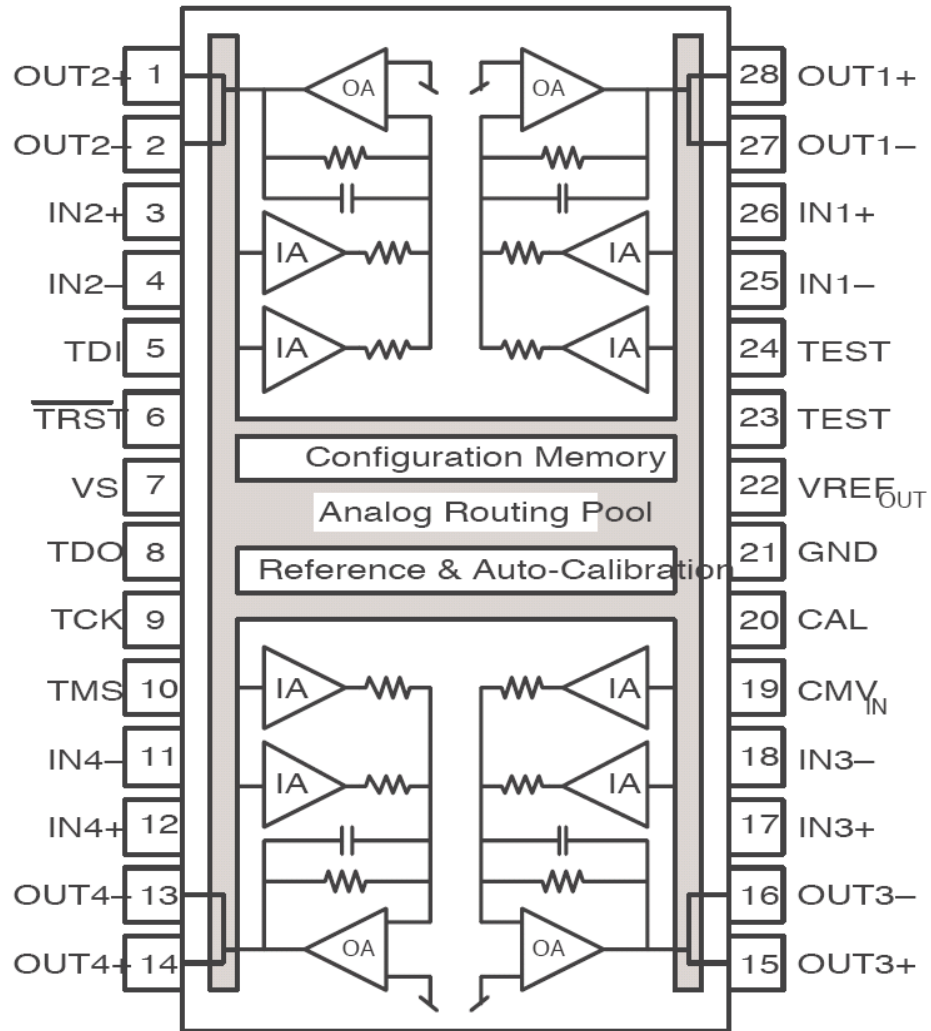
/* Finally send the new bitstream to the
** Virtex Chip */
result = board.setConfiguration(0,bs);

/* And cleanup... */
board.disconnect();

```

Figure 2. Source code for JBits example

from G. Hollingworth et al., "Safe Intrinsic Evolution of Virtex Devices", *Proc. 2000 NASA/DOD Evolvable Hardware Conf.*, 2000



The analog counterpart to the FPGA is the *field programmable analog array* (FPAA). Shown above is the Lattice Semiconductor ispPAC10


The user can program

- input amplifiers gain (by choosing 1 of 8 resistor values plus polarity)
- output amplifier bandwidth (by choosing 1 of 128 capacitor values)
- intrablock and interblock interconnections

The datastream for the ispPAC10 is ≈ 320 bits long. The EA manipulates the binary string directly.

The lowest granularity device is the $0.5\mu\text{m}$ CMOS *field programmable transistor array* (FPTA). This is a prototype chip developed for NASA's JPL.

The FPTA is organized as a 2D array of cells



. Schematic of an FPTA cell consisting of 8 transistors and 24 switches.


Cells can be combined to form current mirrors, logic gates and op amps.

From A. Stoica et al., "Reconfigurable VLSI Architectures for Evolvable Hardware: From Experimental Field Programmable Transistor Arrays to Evolution-Oriented Chips", *IEEE Trans. on VLSI Sys.*, Vol 9(1), 2001

To illustrate how EHW can be used for fault recovery, we will consider an analog FTS. We assume

1. the system is linear with dynamics governed by constant coefficient ordinary differential equations
2. the system is a “black box”—i.e., no access to its inside to repair or replace failed components
3. evolution is done intrinsically

NOTE: black box \Rightarrow compensation is the only means of handling faults

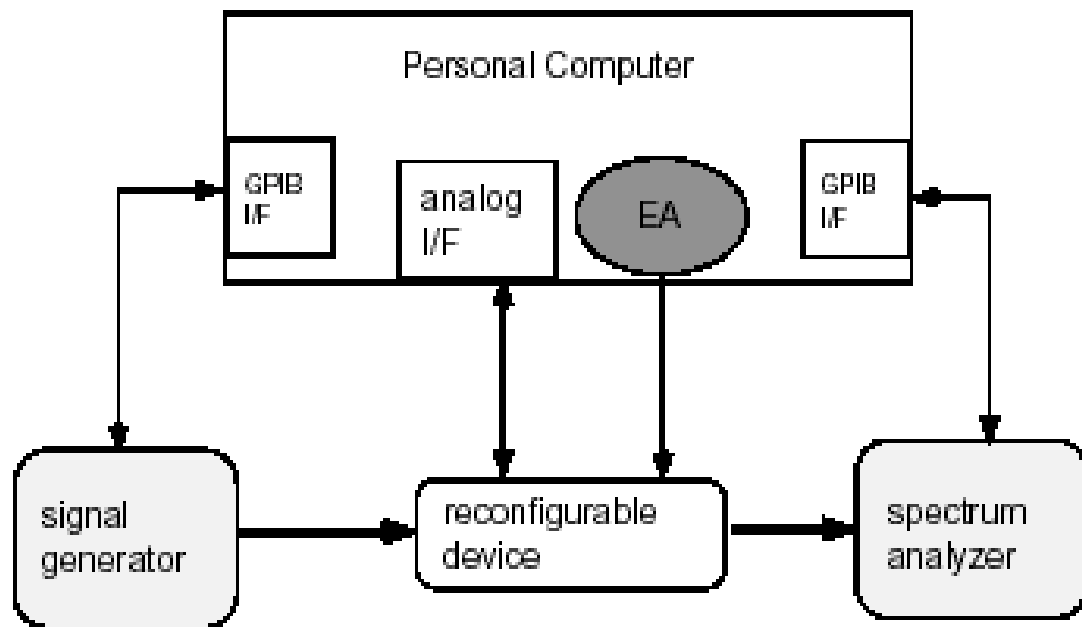


Fault-tolerant analog system

The plant is 3rd order. The fault is manifested by a change in bandwidth; fault recovery should restore the bandwidth.

The reconfigurable device is the ispPAC10, which implements a lead or lag compensator.

The EA evolved a population of 20 for 200 generations using recombination and mutation.



The intrinsic EHW testbench

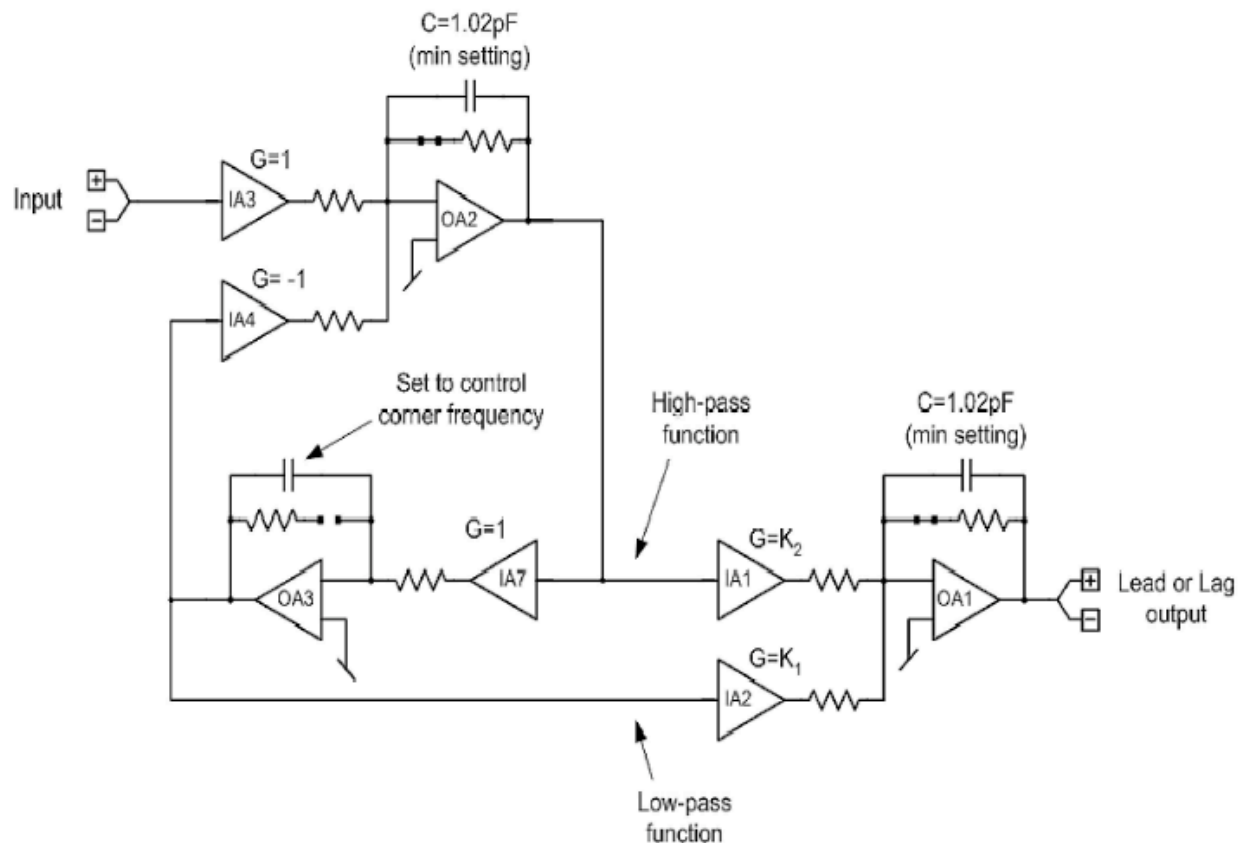
The fitness of configuration C is given by

$$\text{fitness}(C) = \sum_{i=1}^5 [M(i) - M^*(i)]^2$$

where $M(i)$ is the compensated systems magnitude and $M^*(i)$ the desired magnitude at the $i - th$ test frequency.


The 5 test frequencies were chosen such that some were in the passband, one at the -3db point, and some in the stopband.

NOTE: just using the -3db point would produce the trivial solution of just increasing or decreasing the open-loop gain.



The EA had to evolve the capacitor value and the two amplifier gain values (K_1 and K_2).

We also had to evolve two switch positions because the only way to get K_1 or K_2 equal to 0 was to remove the amplifiers IA_1 and IA_2 .



Overview of real-time systems

def (real-time system)

any system that is both logically and temporally correct

def (logically correct)

satisfies all functional specifications

def (temporally correct)

completes all tasks within specified timeframes

Fast does not mean real-time

and

Real-time does not mean fast

For example, consider two real-time delivery systems: one is a courier who guaranties delivery in less than 3 days and an e-mail system that guarantees delivery in 10 minutes.

Notice both systems are logically correct (they guarantee delivery)

But whether or not they are temporally correct depends on the required delivery time.

scenario	delivery	courier	e-mail
1	5 days	X	X
2	5 hours		X
3	5 min		

So, why is RT an issue in fault-tolerance??

ANS: because faults cannot be left uncorrected indefinitely

This impacts intrinsic EHW used for fault recovery because reconfiguration takes time.

In fact, in some cases intrinsic reconfiguration may not even be practical!!!

With intrinsic reconfiguration every solution must be downloaded into the reconfigurable device, which takes time t_p

Device	Type	Size	t_p (ms)
ispPAC10	FPAA	4	100
AN220E04	FPAA	4	3.8
XC3020A	FPGA	64	1.5
Virtex XCV50	FPGA	1728	7
XC4085XL	FPGA	3136	192
APEX II EP2A70	FPGA	6720	12.5
JPL's FPTA2	FPTA	64	0.008

Let λ be the number of new configurations created each generation

Let t_f be the time to conduct a fitness test

Then an EA running for k generations has an intrinsic reconfiguration time of

$$T_r(k, \lambda) = k\lambda(t_p + t_f)$$

But the real problem is t_f and especially for analog systems.

Example:

An AN220E04 FPAA is used to compensate for aging effects in a control system responsible for positioning a satellite's communications antenna.

The reconfiguration search is done by a generational GA run for 500 generations with a population size of 100.

The system's step response is measured to determine if the compensation is correct. This step response test takes $t_f = 625$ milliseconds to conduct.

Hence, $\lambda = 100$, $k = 500$ and $t_p = 3.8$ ms, which makes $T_r(500, 100) \approx 8.7$ hours.

Is 8.7 hours too long??

ANS: maybe...

Reconfiguration times are meaningless unless they are put into context.

For instance, suppose the satellite needs to communicate for a few minutes every 20 hours, and failure to operate can lead to the loss of the satellite.

If a control system fault occurs just after one of these operational periods, no problem.

If the operational period starts in 10 minutes, big problem.

This means you can only determine if a reconfiguration time is too long by comparing it against the fault recovery time.

Since fault recovery has a deadline, EHW-based recovery is a real-time process.

This has strong implications for the EHW community

It is no longer sufficient to just talk about how an EA was able to restore a circuit's functions. (that only shows logical correctness)

Just reporting an algorithm's running time doesn't say anything about temporal correctness either.

In other words,

No EHW-based recovery method can legitimately proclaim efficacy until it is proven to be both logically and temporally correct

Logical correctness is easy to prove. (Try it out and see if it works!)

Temporal correctness is proven by comparing the EA running time against the fault recovery deadline (defined when the FTA or FMEA was conducted).

If both are satisfied, your EHW recovery method is viable.

HOWEVER,

Don't run the EHW algorithm and then compare it against the recovery time.

This is backwards.

Start with the recovery time and then figure out how many generations (G_{\max}) you can run.

Whatever your EA is going to do, its going to have to do it within G_{\max} generations.

Some final comments...

- extrinsic reconfiguration is often impractical as a fault recovery method because
 - The precise nature of the failure may not be known. Hence, any simulation is likely to be inaccurate.
 - It may take too long.

For instance, consider trying to extrinsically reconfigure hardware in a deep space probe operating near the planet Neptune. Communications with the probe will take hours.
- Genetic programming should be avoided for any fault recovery methods. It simply requires far too much computational effort and is unlikely to meet any realistic fault recovery deadline.

More detailed information can be found in

G. W. Greenwood, "On the Practicality of Using Intrinsic Reconfiguration for Fault Recovery", *IEEE Transactions on Evolutionary Computation* 9(4), 398-405, 2005