

Enabling the On-line Intrinsic Evolution of Analog Controllers

David A. Gwaltney
NASA Marshall Space Flight Center
Huntsville, AL 35812
david.a.gwaltney@nasa.gov

Michael I. Ferguson
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
michael.i.ferguson@jpl.nasa.gov

Abstract

The intrinsic evolution of analog controllers to provide closed-loop control of the speed of a DC motor has been previously demonstrated at NASA Marshall Space Flight Center. A side effect of the evolutionary process is that during evolution there are necessarily poor configurations to be evaluated that could cause damage to the plant. This paper concerns the development and implementation of a safe Evolvable Analog Controller (EAC) architecture able to evolve controllers on-line even in the presence of these poor configurations. The EAC concept is discussed and experimental results are presented that show the feasibility of the approach. This EAC architecture represents the first in a series of steps required to make deployment of an evolvable controller a reality.

1 Introduction

Evolvable hardware provides the capability to evolve analog circuits to produce amplification and filter functions. Conventional analog controller designs employ these same functions. Hardware evolution can enable the deployment of a self-configurable analog controller able to adapt to environmental conditions that would otherwise degrade performance, such as temperature varying to extremes or ionizing radiation. Hardware evolution has been widely reported to provide fault-tolerant capability by re-routing internal connections around damaged components or by reuse of degraded components in novel designs. This enables autonomous configuration and recovery from unanticipated conditions to become a reality, making an evolvable controller attractive for use in a remotely located platform, such as a spacecraft. In interplanetary systems, if the speed of self-recovery can be achieved in minutes, extended loss of functionality due to round trip communication times of 40 minutes or more would no longer exist. A spacecraft with self-configuring evolvable controllers could reconfigure to handle control system faults or other unanticipated system changes in the

time it takes for the spacecraft to alert its handlers on the Earth of a problem.

The intrinsic evolution of analog controllers to provide closed-loop control of the speed of a DC motor has been demonstrated at NASA Marshall Space Flight Center [1, 2]. The JPL developed Stand Alone Board Level Evolvable (SABLE) System [3] was the platform for conducting hardware evolution of controller designs. A diagram of the experimental configuration of the initial evolvable controller for motor speed control is depicted in Figure 1. In this case, the reconfigurable analog array is a JPL developed second-generation Field Programmable Transistor Array (FPTA2). In [2] it was demonstrated that the evolvable controller could accommodate simulated internal faults in the reconfigurable analog array by evolving a new controller that could use remaining non-faulty resources. However, a side effect of the evolutionary process is that during evolution there are necessarily poor configurations to be evaluated which could cause damage to the plant. Some of the individuals perform poorly and result in extremely undesirable responses. In a deployed system, this would require a method to disengage the controller from the actuator during evolution. A redundant controller would need to be employed, while the "damaged" controller is being repaired, or a new controller to handle unanticipated changes in the plant is being evolved.

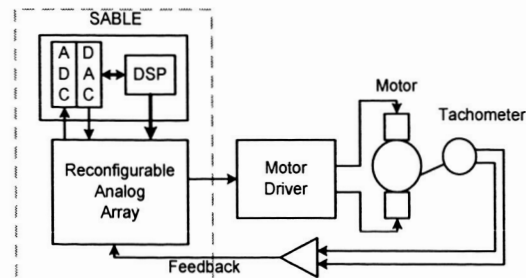


Figure 1 Experimental configuration for the evolution of a motor speed controller

The use of alternative controller architecture can relieve this problem and allow on-line hardware evolution to be used to accommodate faults in the controller and in the plant, or changing plant dynamics.

This paper is concerned with the development and implementation of a *safe* Evolvable Analog Controller (EAC) architecture able to evolve controllers *on-line* even in the presence of these poor configurations. This will be done by using a controller structure that utilizes plant input-output data collected during plant operation. The goal is to produce no undesirable plant response, by evolving new controllers using evolved models of plant dynamics for evaluation. This approach to *safe* hardware evolution of controllers is something that has not been reported in other research efforts concerning controller evolution. This EAC architecture represents the first in a series of steps required to make deployment of an evolvable controller a reality.

2 Approach

The JPL developed Stand-Alone Board Level Evolvable (SABLE) System [3] is used for evolving analog dynamic system models and control electronics. This system employs the JPL designed FPTA2 containing 64 programmable cells on which an electronic design can be implemented by closing internal switches. The schematic diagram of one cell is given in Figure 2. CMOS switches are used to make connections between transistors, variable resistors and capacitors. Each cell has inputs and outputs connected to external pins or the outputs of neighboring cells. More detail on the FPTA2 architecture is found in [4].




Figure 2 FPTA2 Cell Schematic

A diagram of the SABLE system is shown in Figure 1. The main components of the system are a floating-point Digital Signal Processor (DSP), a 100kSa/sec 16-channel DAC and ADC and the FPTA2. There is a 32-bit digital I/O interface connecting the DSP to the FPTA2. The genetic

algorithm (GA) running on the DSP follows a simple algorithm: download an individual, stimulate the circuit with a control signal, record the response, and evaluate the response. The resulting fitness values are used to select an elite percentage of the population and proportional selection is used to select the remaining individuals for the next generation. Then crossover and mutation operators are performed on all but the elite percentage of individuals.

The experimental configuration used to evolve the forward analog transfer function of the controlled plant is depicted in Figure 3. The output of the plant and the evolved transfer function are compared to provide a fitness measure. Initially, the input to the plant is specifically designed to provide excitation that will allow an open-loop transfer function to be identified. A plant inverse transfer function can be identified by making the input to the FPTA2 the output of the plant, and then comparing the output from the FPTA2 with the input to the plant. Both the forward and inverse transfer functions can be used in the implementation of a *safe* EAC architecture, referred to as the EAC in the remainder of the paper.




Figure 3 Experimental configuration for the evolution of a plant transfer function

The EAC is a controller structure that utilizes plant input-output data collected on-line. The controller will produce no undesirable plant response, because new controllers will be evolved in the "background" using evolved models of plant dynamics for evaluation. Because hardware evolution uses high-level design requirements to achieve performance goals using the reconfigurable hardware available to it, evolving analog transfer functions provides the freedom to combine electronic components in any configuration necessary. This will be an advantage in the identification of non-linear transfer functions produced in the presence of system faults. The identified dynamic model will be used in parallel with the plant to evolve new controllers on-line without disturbing the operation of an existing evolved controller, or a stabilizing controller, and its interaction with the plant. The EAC will enable the on-line configuration of controllers without the need for analytical models and will allow the accommodation

of faults, via on-line reconfiguration of the controller structure, when required. In this context, faults include unexpected changes in the dynamics of the plant, in addition to failures experienced by the electronics. When the evolution of a controller has converged or an intermediate controller with desirable fitness is created, the new controller can be switched in to replace the old one, as illustrated in Figure 4. This eliminates deleterious effects on the controller performance due to the evaluation of poor individuals in the population during evolution.




Figure 4 Conceptual EAC architecture

Experimental testing indicates the FPTA2 and SABLE system performs best as a platform for implementing dynamic systems models when the frequencies of interest are between 10 and 10 KHz. Because the useful dynamic range of the vast majority of standard servomotor driven systems is less than 10 Hz, a simulated analog plant is used (Shown in Figure 5). This plant provides a dynamic response similar to a high speed servo that would be encountered in a disc drive application, or another miniature low inertia, high-speed application.




Figure 5 Simulated Analog Plant

The transfer function of the simulated plant is given in Equation 1. The transfer function indicates the simulated plant has a bandwidth of 482 Hz and a steady state gain of 2.5.

$$\frac{V_o}{V_i} = \frac{7575}{s + 3030} \quad (1)$$

3 Evolution of transfer functions

Initially, the evolution of transfer functions is performed using open-loop input-output data from the plant. Just as it is required for standard approaches to linear system identification, excitation signals are selected to excite the salient dynamic response of the plant. An added sine stimulus and a sine sweep stimulus are a good means of providing the input frequency content needed. One type of sine stimulus has advantage over the other at different phases in the evolution of the plant transfer function. Performing evolution in two stages was found to be a profitable approach to system identification of the plant considered in this work.

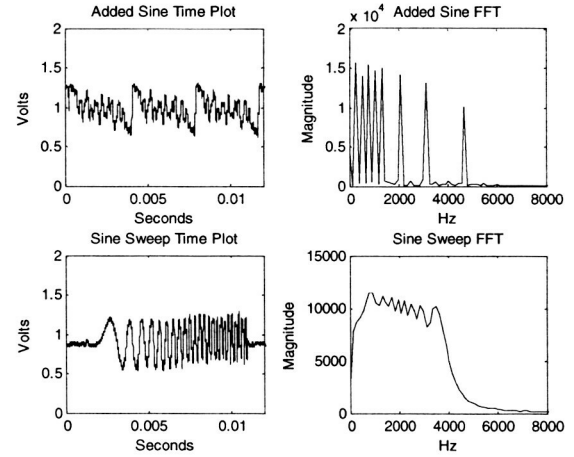


Figure 6 Stimulus signal time and frequency response

Initially, an added sine is used, because it provides a “spiked” frequency response (see Figure 6). When starting from scratch, this provides more information to the genetic algorithm via the fitness evaluation. It was determined that in the initial stage, a fitness based primarily on the differences in the frequency response would provide a better solution. The time response data is used only to indicate the presence of a DC bias between the plant output and the evolved transfer function output. Additionally, no penalty was introduced if the evolved transfer function produces an inverted output. This can be corrected later with an external inverter stage. The fitness function used in the initial stage is represented by Equation 1,

$$F = \left| \frac{P}{n} \sum_{i=1}^n Etime_i \right| + \left| \sum_{j=1}^k M_j * Emag_j \right| \quad (1)$$

where E_{time_i} is the error at the i^{th} data point between the time response of the plant output and the evolved model output, P is a penalty value for the mean of the time response error, n is the total number of time response data points, E_{mag_j} is the error at the j^{th} frequency between the magnitudes of the plant and evolved model frequency response obtained via a Fast Fourier Transform (FFT), M_j is a penalty value for the j^{th} magnitude error, and k is the total number of frequency points for the frequency response.

After a reasonable transfer function is obtained during the first stage, it is refined in the second stage. This refinement includes adding an external inverter, if required, and introducing terms into the fitness based on the time response error. The fitness function for the second stage is given in Equation 2,

$$F = \sum_{i=1}^n |Q_i * E_{time_i}| + \left| \frac{P}{n} \sum_{i=1}^n E_{time_i} \right| + \left| \sum_{j=1}^k M_j * E_{mag_j} \right| \quad (2)$$

where the difference between Equation 1 and Equation 2 is the addition of term that includes the summation of the time response error at the i^{th} data point multiplied by a penalty term Q_i .

In the case of forward transfer functions, it was found that using a sinusoidal sweep (see Figure 6) during the second stage of evolution was beneficial in providing an improved transfer function. A sinusoidal sweep provides continuous coverage of the frequency range of interest.

The physical configuration of the FPTA and connections to it are also of importance. Some experimentation was performed to determine the number of cells needed and what external connections provided benefit. Besides the connection of the excitation signals to the FPTA, other external connections between cells are utilized, but kept to a minimum.

3.1 Forward transfer function

This section presents the best forward transfer function evolved at the time of writing. It was experimentally determined that a transfer function for the simulated analog plant in Figure 5 could be evolved using five cells on the FPTA2 with physical connections as shown in Figure 7. The input and output to the FPTA2 are buffered with unity gain amplifiers so that the evolved transfer function can later be connected to other cells without affecting the response of the evolved circuitry. It is important to include these buffer amplifiers in the configuration during evolution, because the evolutionary process may make use of the input or output impedance of any attached circuitry. Changing the circuitry on the input or output may result in the failure of the evolved circuit. For this evolution cells 4–7 and 16 are active.




Figure 7 FPTA2 Cell physical configuration for evolution of forward transfer function

The initial stage evolution was executed using a randomly initialized population of 200 and was allowed to run for 24,000 generations. In this stage the fitness calculation was primarily based on the differences in FFT magnitude, with the only contribution from the time response being an error measure between the mean values of the plant output and the FPTA output (refer to Equation 1). The fitness reached its minimum (best) value at 18,000 generations, but the evolution ran unsupervised past this point. The resulting evolved configuration for the switches of the five cells resulted in an inverted output that provided the proper frequency response otherwise, as shown in Figure 8. In the time and FFT plots, the output of the plant is drawn in black, while the output of the FPTA2 is shown in gray. Because the FFT plots are very close in appearance, a plot of the error between the FFT magnitudes is also given. Note the FFT Magnitude error is small considering the largest FFT magnitude is on the order of 25,000.

When the output of the FPTA was inverted, there was an offset error. This is visible in the time response error plot in Figure 9, which again includes time and frequency response plots for a sinusoidal sweep stimulus. Note that in the FFT magnitude error plot, the error scale has increased and the differences are more noticeable in the FFT magnitude plot. In the time and FFT plots, the output of the plant is drawn in black, while the output of the FPTA2 is shown in gray.

The second stage evolution was executed using a population of 200 in which one quarter of the population was initialized with the evolved cell internal configuration from the initial stage, while the remaining members of the population were randomly selected. The fitness calculation in this case was dependent on errors between the time and frequency domain responses of the plant and the FPTA2 (refer to Equation 2). The evolution was allowed to run for 3,500 generations.




Figure 8 Time and frequency response comparison after initial stage evolution of forward transfer function

The resulting evolved forward transfer function provides time and frequency responses that are very close to that of the plant. Figure 10 shows both the time response and the frequency responses in which the output of the plant is drawn in black, while the output of the FPTA2 is shown in gray. Note that there is no constant offset in the time response, and the error between the FFT magnitudes has decreased significantly except for the error at DC. This evolved transfer function was used in the evolution of an analog controller for the plant in Figure 5. The experimental results are presented in Section 4.1.




Figure 9 Time and frequency response comparison with sine sweep input and inverted initial stage evolved transfer function output




Figure 10 Time and frequency response comparison with sine sweep input and inverted second stage evolved transfer function output

3.2 Inverse transfer function

The inverse of the plant transfer function is evolved in a similar fashion. The plant inverse transfer function can be identified by making the input to the FPTA2 the output of the plant, and then comparing the output from the FPTA2 with the input to the plant. Finding an inverse transfer function is a more difficult problem, even for simple dynamic systems, so three evolution stages were used. It was experimentally determined that the inverse transfer function can be evolved using eight FPTA2 cells with the external connections shown in Figure 11. The initial stage evolution was executed using a randomly initialized population of 200 and was allowed to run for 16,700 generations. In this stage, the fitness calculation was primarily based on the differences in FFT magnitude, with the only contribution from the time response being an error measure between the mean values of the plant output and the FPTA output (refer to Equation 1).




Figure 11 Cell physical configuration for evolution of inverse transfer function

An added sine input consisting of sinusoids at five different frequencies was used as input to the plant. The resulting evolved configuration for the switches of the eight cells provides an inverted output with a frequency response reasonably close to that of the added sine stimulus used as input to the plant.

The evolved cell internal configuration from the initial stage was used to seed the population for the second stage evolution. Once again this configuration was used to seed one-quarter of the population, while the remaining members were randomly generated to achieve a population size of 200. In the second stage, the output of the FPTA2 was inverted and a slightly modified version of the added sine stimulus depicted in Figure 6 was used. The modification consisted of padding the beginning and the end with “zeros” (0.9V). This was found to make the calculation of the mean more accurate, and aided in generating signals with uniform FFT calculations in each fitness evaluation. Fitness was calculated using both time and frequency response comparison (refer to Equation 2). The fitness reached its minimum (best) value at 17,000 generations, but the evolution ran unsupervised for 21,300 generations. The resulting evolved configuration for the switches of the eight cells produce an output that provides a close approximation to the desired frequency response, as shown in Figure 12. In the time and FFT plots, the output of the plant is drawn in black, while the output of the FPTA2 is shown in gray. Note the time response error and FFT magnitude error is larger in this case than in the forward transfer function case. This is an indicator of the difficulty of identifying inverse dynamics. When a sinusoidal sweep is used as the stimulus and desired response with this evolved configuration, it is seen in Figure 13 that the FFT magnitude errors are larger. Also, the time response is attenuated at higher frequencies, indicated by the increase in time response error as the frequency of the sweep increases.




Figure 12 Time and frequency response comparison with added sine input for second stage evolved inverse transfer function

A third stage evolution was then conducted, using a population of 200, with one-quarter seeded with second stage evolved configuration and the remaining individuals randomly generated. The sine sweep was used as stimulus and the fitness function used both time and frequency response comparison. The evolution was run for 23,600 generations, and the response of the resulting evolved configuration is shown in Figure 14. While the FFT magnitude is closer to the desired response in general, there is now a spike at a low frequency near DC. This manifests itself in the response at the first cycle of the sine sweep. The response is considered more desirable in this case as the generally lower time response error indicates no offset, and will be used as part of a control scheme in the next section.

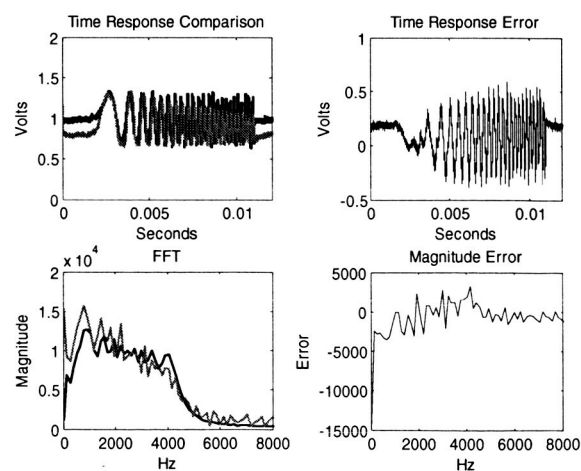



Figure 13 Time and frequency response comparison with sine sweep for second stage evolved inverse transfer function



4 Controller design and response

4.1 Evolution using the forward transfer function

cell 21. Cell 21 must provide the difference function for the controller, as well as the dynamics. Cell 20 is primarily used to provide biasing for cell 21.

Figure 15 Diagram of cell physical configuration for evolution of an analog controller

After connecting the evolved controller to the analog plant, the reference signal was varied in shape, frequency and amplitude. Figure 16 shows time response plots for four different reference signals. The reference signal is black, while the plant output is gray. In each case, the controller was acceptable at controlling the plant, with small offset between the reference and the plant output.




Figure 16 Time response plots for control of the analog plant of Figure 5 using the evolved controller

4.2 Direct application of inverse transfer function

The evolved inverse transfer function was directly applied to the control of the analog plant in Figure 5. The experimental configuration is shown in Figure 17. In this case, the control is open-loop, as there is no feedback path from the plant output to the controller. If the inverse transfer function has been precisely identified, the plant output will closely follow the reference input.




Figure 17 Experimental configuration for open-loop control of the plant using the inverse transfer function

Figure 18 shows four time response plots from experimental runs for different reference inputs. The reference signal is black, while the plant output is gray. As can be easily seen, the evolved inverse transfer function does not precisely match the inverse of the analog plant. While the plant output closely follows the square wave in the upper left plot, there is a large difference between a reference input of a sinusoid and the resulting plant output as shown in the upper right hand plot. The lower left hand plot is the response to a higher frequency square wave, with a different mean value (offset). The rounded corner of the rising edge of the plant response is due to the increased frequency of the reference signal. But the large offset between the reference and the plant output is due to the change in the mean value of the reference input. The ramp function in the lower right hand corner clearly shows the shortcomings of the evolved inverse transfer function. The ramp signal sweeps across most of the dynamic range of the

plant, and the plant response should be much more linear, if the inverse model was closer to the actual plant inverse dynamics. More work needs to be done to evolve a better inverse transfer function, but what was obtained is in the so-called “ball park.” In practice, the inverse model will be used as a feed forward control component along with a simple feedback controller. In this way, the feedback controller can help to null out offsets and other errors due to inaccuracies in the inverse model. This architecture will be explored further in future work.




Figure 18 Time response plots for control of the analog plant of Figure 5 using the evolved inverse transfer function

5 Summary

The work presented here is considered part of a series of steps that will be required to implement an EAC that can evolve to provide on-line reconfiguration in a deployed control system. The significant issue of protecting the controlled plant from the evaluation of poor individuals can be addressed by the use of evolved transfer functions in the evolution of a controller. The feasibility of this approach has been shown in the experimental results presented here.

The advantage of using evolved transfer functions is that they can be adapted to represent un-expected dynamics or non-linearities due to failures that arise in a plant, without the need for analytical models. Evolution requires no analytical knowledge and can be used to autonomously identify plant transfer functions. While simulation approaches exist for representing such systems, they generally require significant a-priori analysis by a human designer, or may require extensive computing capability. In a remote system, such as a spacecraft, neither a human designer nor extensive computing resources are available. The approach presented here can provide simulated transfer

functions in analog circuitry which can accurately approximate continuous time systems.

It is desirable that the plant transfer functions be identified during normal operation of the controlled plant. In this work, special stimulus was applied to an open-loop plant for use in identification. While this may be initially required to get a baseline transfer function that can be modified, ultimately the transfer function should be continuously updated to track changes in the plant during normal operation. Then, the EAC will be prepared to modify the controller quickly, when the need arises.

Finally, the EAC must be proven with more complex plant dynamics. The simple plant dynamics of the analog plant used here are good for developing concepts and approaches. The next step is to increase the complexity of the controlled plant. First a second order system with resonance (complex poles) will be used for the plant, with the next goal being a fourth order plant. The refinement of the approach presented here is one aspect of a research proposal submitted by the authors in response to NASA Research Announcement NRA2-38169. This proposal has been selected for award, and the research and development activities commenced in March 2005. Other aspects of the proposed work include the use of domain knowledge in the evolutionary algorithm and the design of an evolution oriented reconfigurable device specifically targeted for control applications. The FPTA2 represents an important innovation in the creation of evolution oriented reconfigurable electronics. However, it was not designed for analog control applications, and experimental work using in the FPTA2 in control applications has identified the need for a reconfigurable device with different capabilities. A new chip design targeted for control applications will facilitate efficient evolution of control structures and provide the necessary components to support plants with lower bandwidth.

References

- [1] Gwaltney, D. A. and Ferguson, M. I., *Hardware Evolution of Analog Speed Controllers for a DC Motor*, E. Cantu-Paz, et. al., (Eds.), GECCO 2003, LNCS 2723, July 2003, pp. 442-453.
- [2] Gwaltney, D. A. and Ferguson, M. I., *Intrinsic Hardware Evolution for the Design and Reconfiguration of Analog Speed Controllers for a DC Motor*, 2003 NASA/DoD Conf. On Evolvable Hardware, Chicago, IL, July 2003.
- [3] Ferguson, M. I., Zebulum, R., Keymeulen, D. and Stoica, A., *An Evolvable Hardware Platform Based on DSP and FPTA*, Late Breaking Papers at the Genetic and Evolutionary Computation Conference (GECCO-2002), July 2002, pp. 145-152.

- [4] Stoica, A., Zebulum, R., Keymeulen, D., *Progress and Challenges in Building Evolvable Devices*, Evolvable Hardware, Proceedings of the third NASA/DoD Workshop on, July 2001, pp 33 – 35.