Proceedings of the 9th European Conference on
Artificial Life (ECAL 2007), Lisbon, Portugal,
September 2007.

Evolution of an Adaptive Sleep Response in Digital Organisms

Benjamin E. Beckmann, Philip K. McKinley, and Charles Ofria

Department of Computer Science and Engineering
Michigan State University
East Lansing, MI 48824, USA
{beckma24,mckinley,ofria}@msu.edu

Abstract. Adaptive responses to resource availability are common in natural systems. In this paper we explore one possible evolutionary cause of adaptive sleep/wake behavior. We subjected populations of digital organisms to an environment with a slowly diminishing resource and recorded their ability to adapt to the changing environment using sleep. We also quantified the selective pressure not to sleep in this competitive environment. We observed that diminishing resource availability can promote adaptive sleep responses in digital organisms even when there is an opportunity cost associated with sleeping.

Keywords: Digital evolution, digital organism, Avida, adaptive behavior, sleep, resource availability.

1 Introduction

A population of organisms in an environment where a resource is always available can be non-adaptive and function exceptionally well. There is little or no selective pressure on the organisms to adjust their behavior within this environment since resources are plentiful and can be consumed at any time [1]. If resources often become diminished or unavailable, an adaptive response might allow for more conservative resource usage [2] or increased energy storage [3]. Natural organisms often display adaptive behavior that coincides with environmental changes where resources fluctuate [4, 5]. An example of this type of adaptive response occurs in nocturnal rodents and insects that sleep during the day and forage for food under the cover of darkness. Animals that hibernate also display an adaptability that allows them to avoid extended periods of low resource availability by increasing the size of their fat stores prior to hibernation [6].

This form of adaptive behavior in natural organisms serves multiple purposes. During sleep periods an animal rests [7], reprograms its brain [8] and performs internal maintenance tasks [9]. However, while an animal is in a state of slumber it is less aware of its environment. How could resource-aware adaptive behaviors, such as sleep and hibernation, have evolved in competitive environments where torpid organisms are vulnerable to active organisms? Is there a selective pressure to sleep caused by resource limitations in environments with periodic resource availability? The remainder of this paper attempts to answer these questions through experiments with digital organisms.

Previous work has been done in this area using neural networks [10]. In [10], the organisms were subjected to two different environments with periodic light availability, where the organism's ability to find a resource was impaired relative to the current light intensity. It was shown that the combination of a biological clock and light sensor produced the best results in an environment where light readings may not correctly disambiguate day from night. The work presented in this paper differs for [10] in that it does not impose a predefined structure on the organisms, provide a common starting point to the organisms, or give any information, ambiguous or not, to the organisms directly. All of these mechanisms must be evolved while preserving an organism's ability to self-replicate and while avoiding other detrimental behavioral changes. We begin with a brief overview of the Avida digital evolution platform [11] and the experimental setup, followed by presentation of the experimental results.

2 Avida

Avida is an experimental software platform for research in computational evolutionary biology [11]. In the past several years, Avida has been used to conduct pioneering research on the evolution of biocomplexity [12–14]. Avida provides researchers with tools to study the evolutionary process in greater detail and less time than previously possible.

In an Avida experiment, self-replicating digital organisms compete against each other in a fixed-size steady-state population. As shown in Fig. 1, each organism resides in a *cell* (one organism per cell) and comprises a circular list of assembly-like instructions (its genome) and a virtual CPU capable of executing those instructions. Cells are organized according to a topology; in this study we used a two-dimensional grid. Every virtual CPU has three 32-bit registers (AX, BX, and CX) and two stacks capable of storing up to ten 32-bit numbers. The virtual CPU has an instruction pointer (IP) that determines which instruction in an organism's genome to execute. The IP can be moved throughout the genome with the use of conditional if-statements and explicit move instructions. While the Avida instruction set is a Turing complete language, only basic computational instructions are available and complex computations must be constructed by combining simple instructions (i.e, nand, inc, and add) with the input/output instruction.

We record all input and output to and from each organism in the population and examine them to determine the computational *tasks* performed. An example task is the bitwise-and of two numbers [13]. To complete this task an organism must read in two numbers and output the bitwise-and of those two numbers sometime in the future. We added an energy model to Avida that allows an organism to obtain and store energy. Once a task has been completed, an *energy* reward is added to the organism's current energy. In this study, the size of the energy reward is subject to the availability of *resources* in the environment.

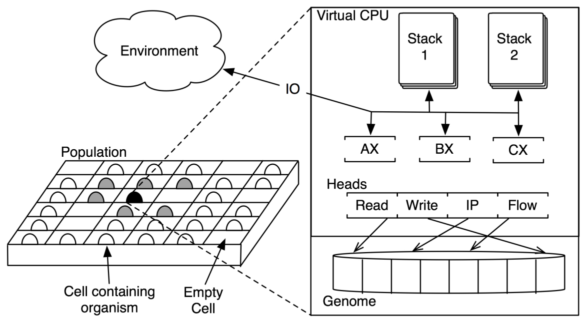

Fig. 1. Avida population and structure of a single organism.

Organisms can sense the quantity of resources within the environment. The more plentiful the resource, the larger the reward for performing a task.

An organism's current energy level is used to determine its metabolic rate, as shown in (1). Avida uses a probabilistic scheduler to assign virtual CPU cycles to organisms in the population. Organisms with higher metabolic rates are assigned higher priority within the scheduler, and therefore execute more instructions relative to organisms with lower metabolic rates. The metabolic rate is inversely proportional to a user-defined variable, $InstructionsBeforeZeroEnergy$, which specifies how many instructions an organism can execute before it runs out of energy, given no new energy influx. Probabilistically, organisms with a higher metabolic rate will execute more instructions and produce more offspring than those with less energy.

$$MetabolicRate = \frac{Energy}{InstructionsBeforeZeroEnergy} \tag{1}$$

Avida organisms are responsible for their own replication through the use of replication-specific instructions. To reproduce, an organism must perform three distinct functions: allocate space at the end of its genome for its offspring's genome, duplicate its own genome instruction by instruction into that space, and divide the resulting genome into two parts. Upon division, the parent organism's state is reset, the parent's energy is divided equally between itself and its offspring, and the offspring's genome is used to create a new organism. The offspring is placed in a random cell in the grid, replacing and terminating any organism that previously occupied that location. Variation among organisms in the population occurs when instructions are copied. Each copied instruction is subject to three types of mutation (modifying the instruction, deleting the instruction, or inserting an additional instruction) that occur at user defined rates. Replication is asexual, and therefore every Avida run presented here begins with the same single organism that serves as an ancestor for all successive organisms in the population. Each run is started with a different random number seed, resulting in different evolutionary paths taken by the population.

3 Experimental Setup

In these experiments, the population of digital organisms is arranged in a 60×60 grid. When an instruction is being copied there is a 0.75% chance that the instruction being copied will be mutated. During replication there is a 5% chance an instruction will be deleted, and a 5% chance that a random instruction will be inserted. On average each organism in the population will execute one instruction per *update*, the standard unit of time in AVIDA.

As in [13], organisms are rewarded for performing tasks that are Boolean logic operations. Specifically, we used the five tasks listed in Table 1. Each task has an associated reward, indicating the number of energy units an organism gains when completed, and a limit on how many times a individual organism may be rewarded for performing it. Completing even these relatively simple tasks can require several instructions. Table 2 shows a "hand-built" solution for the AND task (a NOP instruction modifies the behavior the preceding instruction, for example, placing the result in a different register than the default). Of course, evolution may produce many different solutions for the same task.

The environment contains a single resource that is available periodically. When the resource is available, it is non-depletable, and all five tasks described in Table 1 are maximally rewarded. If an organism completes a task when the resource is unavailable, no reward is given. The duration of the resource availability changes throughout every experiment except the control experiment, where it remains constant. Resource availability is defined in "years" and "days." Each year consists of 500 days, each of which lasts for 256 time steps (updates). During each year, the availability of the resource remains constant. That is, each day of a year has the same duration of resource availability. At the beginning of each day the resource becomes available for a period of time depending on the current year. For the first year the resource is available during 100% of the day. After each passing year, the availability of the resource during a day is reduced by 6.25% of a full day until it becomes zero, which deprives the population of energy and eventually brings on its demise. Through evolutionary change brought upon by depriving the population in this manner, we observe under which conditions the population of digital organisms will find sleep useful.

Table 1. Rewarded tasks.

Task Name	Input	Bitwise Output	Reward	Max Times Rewarded
ECHO	A	A	1000	35
NAND	A, B	$\neg(A \wedge B)$	1500	20
NOT	A	$\neg A$	1500	20
ORNOT	A, B	$A \vee (\neg B)$	2000	13
AND	A, B	$A \wedge B$	2000	13

We have added six instructions to the base AVIDA instruction set, enabling an organism to sense and respond to its environment. These instructions are: TIME, SENSE, and four variations of SLEEP. Executing the TIME instruction stores the

Table 2. Instruction sequence that when executed completes the AND task.

Instruction	AX	BX	CX	Stacks 1,2	Output	Description
IO	?	X	?	?,?	?	read X into BX
IO nop-C	?	X	Y	?,?	?	read Y into CX
nand	?	X nand Y	Y	?,?	–	BX $\leftarrow \neg$(AX $\wedge$ BX)
push	?	X nand Y	Y	X nand Y, ?	–	push BX onto stack 1
pop nop-C	?	X nand Y	X nand Y	?,?	–	pop stack, place result in CX
nand	?	X and Y	X nand Y	?,?	–	BX $= \neg$(BX $\wedge$ CX)
IO	?	Z	X nand Y	?,?	X and Y	output BX

current time step in a register within the organism's virtual CPU. The SENSE instruction allows an organism to detect the presence or absence of the resource; it loads one of the calling organism's registers with the current quantity of the resource times 100. (The value of the resource is multiplied by 100 to allow for a wider range of the sensed value.) The SLEEP instructions allow organisms to enter a low energy state that lasts for multiple CPU cycles. Compared to other instructions, the SLEEP1-4 instructions cost 100 times less energy to execute and last for 10, 20, 40, and 80 times more CPU cycles, respectively.

To help answer the questions posed in Section 1 we ran three experiments. The first is a control where the resource is available 100% of the time. In the second experiment resources are diminished for the duration of each run. In the final experiment the sleep instructions have been replaced by a null instruction to quantify the selective pressures being applied to the SLEEP instructions.

4 Experimental Results and Discussion

We define an environment where a resource is available for the duration of each day. In this environment, which is referred to as "constant," the organisms in the population do not benefit from an adaptive response based on the availability of the resource because the resource can be used at any time. The remainder of this paper presents evidence that a decline in resource availability within a single-resource environment can produce an adaptive resource-aware response.

To test this hypothesis we conducted two experiments; results presented are the average of 50 runs. In the first experiment the resource is available during the entire run (constant environment) and in the second experiment the availability of the resource is reduced over the course of the run (declining environment). Figure 2(a) displays the average metabolic rate in both the constant and declining resource environments. For clarity, error bars are omitted; the maximum standard error is 0.018 for constant environment and 0.01 for declining environment. The 16 vertical lines in Fig. 2(a) denote years, where a 6.25% decrease in resource availability occurs in the declining resource environment. As shown, the metabolic rate in the constant environment tends to stabilize as the run proceeds, but decreases over time in the environment with declining resource availability. This behavior is expected, since organisms can receive rewards for completing tasks continually in the constant environment, but less often as time lapses in the declining resource environment. In fact, after the last vertical line the organisms in the declining resource environment populations no longer have a source of energy, and eventually the populations will die when they run out of stored energy.

Figure 2(b) shows the average maximum and minimum number of organisms sleeping at some time during a day in each environment. The maximum and minimum numbers of organisms sleeping during a day in the constant environment remain relatively close together. In contrast, organisms in the declining resource environment have evolved to participate in inactive periods, where at the peak, on average, greater than 10% of the organisms in the population are sleeping. At this point the number of organisms sleeping in the declining environment is significantly above the number sleeping in the constant experiment. (p-value < 0.0003, using Wilcoxon rank sum test for equal medians). A sample of evolved code from one of the runs is given in Table 3. The code produces a resource-aware behavior when executed. Specifically, the organism enters a loop that ends when the resource becomes available.

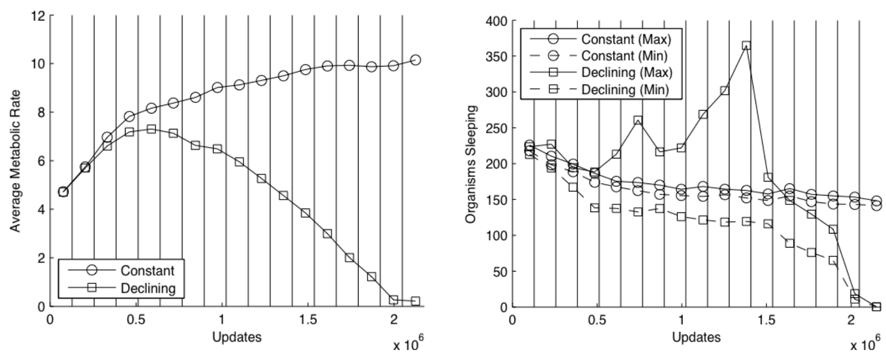

(a) Average metabolic rate (b) Average maximum/minimum sleeping organisms

Fig. 2. Comparison of sleep responses in two environments, one where the resource is available 100% of the time (constant), and one where the resource availability decreases over time (declining). Results are the average of 50 runs.

Since the organisms sleep more in the declining resource environment, one might infer that the organisms accumulate more SLEEP instructions in their genomes. However this is not true. Figure 3(a) shows the number of SLEEP instructions that are *present* in the organisms' genomes in both environments, along with the number of sleep instructions *executed* in each. For the first half of the runs, organisms in both environments have substantially more SLEEP instructions in their genomes than they actually execute. The gap then begins to narrow in the declining environment, and by the end of the runs the number of executions nearly equals the number present. The increase in the execution of SLEEP instructions in this environment suggests that sleeping is more beneficial as the resource availability diminishes. Figure 3(b) shows the rate of execution of

Table 3. Evolved code that loops until the resource becomes available.

Instruction	Explanation
H-SEARCH	place FLOW-HEAD at next instruction
SLEEP	start sleeping
SENSE	read resource availability into BX register
IF-EQU-0	if BX $\neq 0$ skip next instruction
MOV-HEAD	move INSTRUCTION-HEAD to FLOW-HEAD

SLEEP instructions over the course of the runs in the declining resource environment. As expected, the SLEEP instructions with lower CPU cycle costs are used more heavily than the more expensive SLEEP instructions, especially early in the runs. As the resource becomes scarce, the number of more expensive SLEEP instructions increases. This adaptation allows for longer sleep cycles with fewer executed instructions.

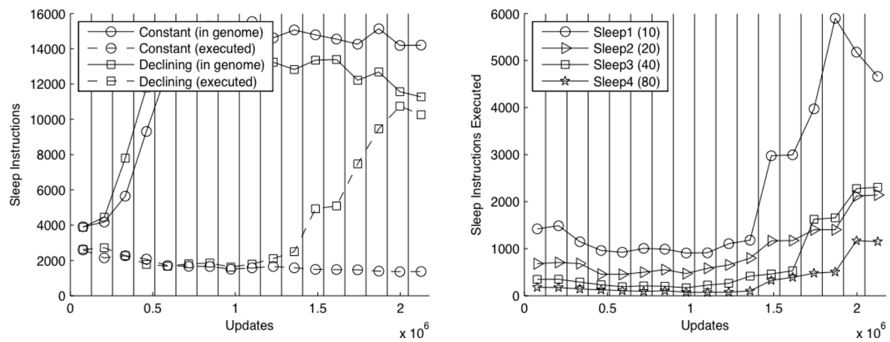

(a) number of SLEEP instructions present in and executed by organisms

(b) number of four SLEEP instructions executed, declining environment.

Fig. 3. Number of SLEEP instructions present in and executed by organisms in the constant and declining environments. Average over 50 runs.

When AVIDA organisms are exposed to an environment where resource availability varies during a day, they evolve an adaptive resource-aware response. An example is shown in Figure 4, which depicts snapshots of the 60×60 grid during a single day in a population that evolved this adaptive sleep/wake behavior. The black squares depict organisms that are sleeping. At this point in the run, the resource is available for the first 112 (out of 256) updates. Figure 4(a) shows the population at the beginning of a day. Figure 4(d) shows the population at the day's midway point where the resource is no longer available and organisms are beginning to enter a sleep cycle. During this day the peak number of organisms sleeping at one time is 2111 or 58.6%, shown in Figure 4(e). After this point the organisms start to wake up and await the next period of resource availability.

Figure 5(a) plots the number of organisms sleeping and the resource availability during three consecutive days near the midpoint of a single run, when the resource is available during the first 50% of each day. As shown, there is a tight correlation between number of sleeping organisms and lack of resources. Examination of evolved genomes shows that organisms in this population have

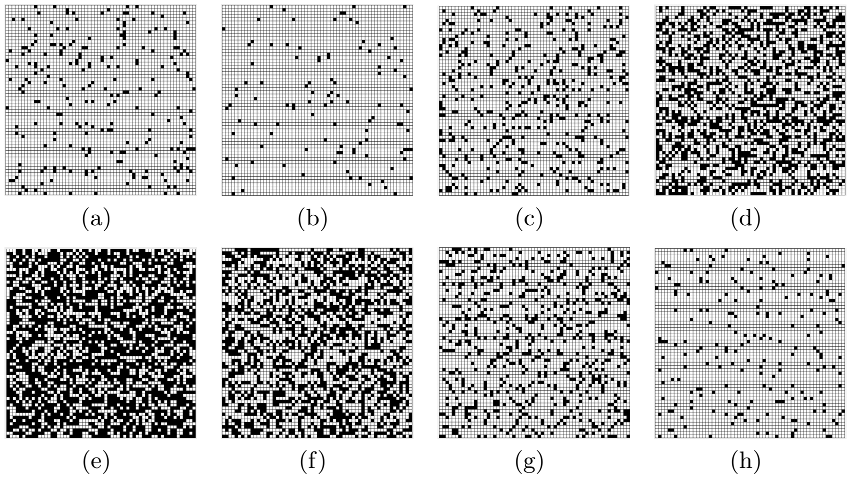

Fig. 4. Representations of a population's response to the resource availability over a single 256 time-step day. Black squares represent sleeping organisms and white squares represent awake organisms. The resource is available for the first 112 time steps. a) t = 1, 231 sleeping, resource becomes available; b) t = 64, 108 sleeping; c) t = 128, 469 sleeping; d) t = 152, 1355 sleeping, resource is no longer available; e) t = 180, 2111 sleeping; f) t = 204, 1502 sleeping, organisms are beginning to wake up; g) t = 228, 667 sleeping; h) t = 256, 189 sleeping, day ends and resource becomes available again.

evolved to begin their sleep cycle just before the beginning of resource deprived periods, and begin preparing data to be used in tasks, just prior to the return of the resource. This "early to bed, early to rise" behavior allows organisms to finish tasks early during periods of resource availability, thereby increasing the probability of receiving a reward. It also helps to avoid situations where an organism's execution is delayed, causing a task to be completed just after the resource disappears, in which case the organism receives no reward. This adaptive behavior arose in 37 out of 50 runs in the declining resource environment.

Although the populations evolved an adaptive behavior, in the above trials the fraction of concurrent sleeping organisms never stabilized above 60%. To help explain why more organisms did not sleep, we conducted a final experiment, where the four SLEEP instructions were replaced by the NOP-X instruction, which has no effect on the virtual CPU when executed, and has CPU and energy costs equal to the non-sleep instructions. The same experimental setup with a declining resource availability was used, the only difference being the replacement of the SLEEP instructions with NOP-X. Figure 5(b) compares the number of SLEEP and NOP-X instructions present and executed in the populations. In both cases the NOP-X instruction is significantly more plentiful than the SLEEP instructions. In fact the p-values for both are less than 0.0001. Selective pressures produced by this treatment favored doing nothing for 1 CPU cycle and paying a higher energy cost, over doing nothing for multiple CPU cycles and using 100 times less energy. Yet, even in the presence of this selective pressure, an adaptive resource-aware sleep/wake behavior has evolved to a point where a majority of the organisms in a single population sleep at the same time.

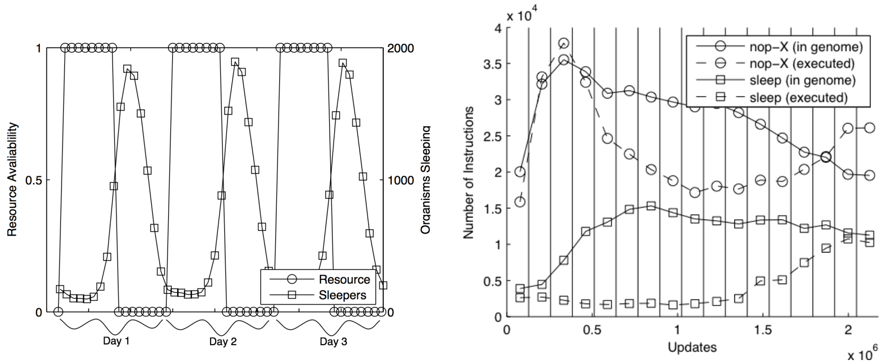

(a) (b)

Fig. 5. (a) Attempted resource usage by organisms (resource activity) and resource availability vs. time for a typical 3-day interval. (b) A comparison of SLEEP instructions (squares) to inert NOP-X instructions (circles); solid lines indicate the frequency with which each instruction is found in the genome and dashed lines indicate the frequency at which they are executed.

5 Conclusion

Revisiting the questions posed in Section 1, we have shown that populations of digital organisms are capable of evolving resource-aware adaptive sleep/wake behavior in an environment where resource availability is periodic and declines over time. The organisms in these populations become highly active when the resource is available and sleep when it is not. This behavior evolves even though sleeping organisms are vulnerable to non-sleepers and there exists a selective pressure not to sleep. This behavior evolved and remained stable in a majority of the populations in our experiments. We also have seen evidence suggesting that the adage "early to bed, early to rise" describes an evolved behavior, as organisms maximize their probability of being rewarded for completing tasks. This behavior evolved even in the presence of a selective pressure not to sleep.

Continuations of this work using additional environments are ongoing. Environments with added costs, instruction and environmental impairments, positive and negative reinforcement, and punishment will all be tested for effectiveness. Additionally, seasonal resource availability models are under development and will be used to model the natural world more closely. Finally, environments encouraging predator/prey relationships will be examined for evidence of coexisting diurnal and nocturnal behaviors among organisms within the same population.

Further Information. Papers on digital evolution and the AVIDA software are available at `http://devolab.cse.msu.edu`. Information on evolving adaptive and cooperative behavior can be found at `http://www.cse.msu.edu/thinktank`.

Acknowledgments. The authors gratefully acknowledge the contributions of the following individuals to this work: David Knoester, Jeffrey Clune, Sherri Goings, David Bryson, Richard Lenski, Heather Goldsby, and Betty Cheng. This work was supported in part by the U.S. Department of the Navy, Office of Naval Research under Grant No. N00014-01-1-0744, National Science Foundation grants EIA-0130724, ITR-0313142, and CCF 0523449, and a Quality Fund Concept grant from Michigan State University.

References

1. P. J. Kennedy and T. R. Osborn, “Operon expression and regulation with spiders,” in *2000 Genetic and Evolutionary Computation Conf. Workshop Program* (D. Whitley, D. Goldberg, and E. Cantu-Paz, eds.), pp. 161–166, 2000.
2. T. Wang, C. C. Hung, and D. J. Randall, “The comparative physiology of food deprivation: From feast to famine,” in *Annual Review of Physiology*, vol. 68, pp. 223–251, Annual Reviews, 2006.
3. M. M. Humphries, D. W. Thomas, and D. L. Kramer, “The role of energy availability in mammalian hibernation: A cost-benefit approach,” in *Physiological and Biochemical Zoology*, vol. 76, pp. 165–179, University of Chicago Press, March-April 2003.
4. R. J. Denver, N. Mirhadi, and M. Phillips, “Adaptive plasticity in amphibian metamorphosis: Response of Scaphiopus Hammondiitadpoles to habitat desiccation,” in *Ecology*, vol. 79, pp. 1859–1872, Ecological Society of America, 1998.
5. G. H. Johnsen and P. J. Jakobsen, “The effect of food limitation on vertical migration in Daphnia Longispina,” in *Limnology and Oceanography*, vol. 32, pp. 873–880, American Society of Limnology and Oceanography, 1987.
6. R. A. Young, “Fat, energy and mammalian survival,” in *American Zoologist*, vol. 16, pp. 699–710, The Society for Integrative and Comparative Biology, 1976.
7. H. Zepelin, “Mammalian sleep,” in *Principles and Practice of Sleep Medicine* (M. Kryger, T. Roth, and W. Dement, eds.), pp. 82–92, Philadelphia: Saunders, 2000.
8. F. Crick and G. Mitchison, “The function of dream sleep,” *Nature*, vol. 304, no. 5922, pp. 111–114, 1983.
9. W. H. Moorcroft, *Sleep, Dreaming, and Sleep Disorders.* MD: University Press of America, Inc., 2nd ed., 1993.
10. M. Mirolli and D. Parisi, “Artificial organisms that sleep,” in *Advances in Artificial Life - Proceedings of the 7th European Conference on Artificial Life* (W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, and J. Ziegler, eds.), vol. 2801 of *Lecture Notes in Computer Science*, (Dortmund, Germany), pp. 377–386, Springer, September 2003.
11. C. Ofria and C. O. Wilke, “Avida: A software platform for research in computational evolutionary biology,” in *Journal of Artificial Life*, vol. 10, pp. 191–229, International Society of Artificial Life (ISAL), March 2004.
12. R. E. Lenski, C. Ofria, T. C. Collier, and C. Adami, “Genome complexity, robustness, and genetic interactions in digital organisms,” in *Nature*, vol. 400, pp. 661–664, 1999.
13. R. E. Lenski, C. Ofria, R. T. Pennock, and C. Adami, “The evolutionary origin of complex features,” in *Nature*, vol. 423, pp. 139–144, 2003.
14. C. Adami, C. Ofria, and T. C. Collier, “Evolution of biological complexity,” *Proceedings of the National Academy of Sciences*, vol. 97, pp. 4463–4468, 2000.