To appear, Dorin, A., "A Co-Evolutionary Epidemiological Model for Artificial Life and Death", in *Proceedings 8th European Conference on Artificial Life*, Capcarrere et al (eds), Springer-Verlag, LNAI, 2005

A Co-Evolutionary Epidemiological Model for Artificial Life and Death

Alan Dorin
Centre for Electronic Media Art, Computer Science & Software Engineering,
Monash University, Clayton, Australia 3800
aland@csse.monash.edu.au

Abstract. This paper presents a model of the co-evolution of transmissible disease and a population of non-randomly mixed susceptible agents. The presence of the disease elements is shown to prevent the onset of genetic convergence of the agent population. The epidemiological model also acts in a distributed fashion to counter the tendency of the agent population to occupy spatially close-knit communities. The simulation applies a modified mathematical SIR epidemiological model of disease transmission in combination with the well-studied technique of artificial ecosystems. It includes various aspects of disease transmission that are not usually modelled due to the effort required to incorporate them into mathematical models. These include a distributed agent population with non-uniform infectiousness and immunity as well as a mutable disease model with evolving latency and infections that evolve to prey on diverse agent characteristics.

Keywords. epidemiological model, co-evolution, artificial death, ecosystem

1 Motivation and Past Work

Digital evolutionary simulations may converge with the population predominantly of similar genetic composition. This convergence is undesirable if it occurs (for example, due to the presence of local maxima in the fitness landscape) before a solution to an optimisation problem has been discovered. Hence where the desired simulation result is an exploration of diverse solutions within some constraints, early convergence needs to be avoided. This is also the case when the problem is to generate an interesting, evolving agent population of sonic or visual forms for an artwork or interactive installation [1]. For this latter reason it was decided to explore a strategy for preventing the convergence of a population that would be applicable generally to agent simulations.

To meet its aim, this paper prefers an elegant, emergent, decentralized approach over that of a hard-coded or centralized controller. Hence, it focuses primarily on a notable omission from many Artificial Life models and publications, *disease*. Typical Artificial Life ecological simulations model creatures competing for food, mating, fighting, and dieing. Yaeger's *PolyWorld* is a seminal example in which agents interact utilizing colour vision [2]. Todd has noted strategies for removing creatures from a population subject to a genetic algorithm but stops short of exploring different reasons for their death [3] (for example disease or suicide). Mascaro et al. have dealt specifically

with suicide in a population of simple agents [4]. Ray's *Tierra* simulation eliminates elderly or ineffective population members with a "reaper". Also of interest is the emergence of "parasitic" code in his system [5]. The usual methods of removing members of a population may also have been employed – culling random agents, unfit agents or replacing parents with their offspring for example. These standard approaches do not improve the diversity of the population at any one time, in fact if carelessly applied they may be responsible for its convergence.

The Artificial Life literature has much to say on co-evolution as a means of improving a genetic algorithm's performance through increased population diversity [6, 7]. This work is similarly inspired, only the simulation models virtual worlds and does not optimise explicit fitness functions. The co-evolutionary model presented here is novel also since the disease/parasite that co-evolves with the agent population is wholly dependent for its position in space (and in fact its existence) on the susceptible agents.

1.1 An Introduction to Models From Epidemiology

A fully-cited history of the mathematical theory of epidemics is beyond the scope of this paper. The history leading to the classic model discussed below is provided in [8, 9].

At least since the 1920's, stochastic models of epidemics have been utilized. The standard model is based on a population of individuals who are either susceptible to a specific disease (*susceptibles* denoted **S**) or infected with the disease and capable of transmitting it to others (*infectives* denoted **I**). Population members who overcome a disease may become immune to further infection[1] or may become susceptible once again depending on the particular disease. Population members who are immune to a disease or remain infected but through isolation cannot transmit it, are considered *removed* (denoted **R**). The model as described is known as an *SIR* model. It may be modified slightly to provide fresh susceptibles through birth or immigration.

Some pertinent parameters of epidemic models are as follows. The period of time during which a disease exists entirely within an organism is known as the disease's *latent period*. The organism is not infective during this period. An *incubation period* often follows latency. During incubation the organism may not show outward sign of infection but is nevertheless infective. Usually once the incubation period is over, the victim of the disease is clearly marked by symptoms and can therefore be avoided by susceptibles.

Probabilistic epidemiological models that operate in discrete time steps are particularly suited to implementation in software.[2] At any time step, the probability of a new case of the disease appearing is proportional to the number of susceptibles multiplied

---

[1] Following a bout of a disease a victim may be deceased, alternatively their immune system may prevent repeat infiltration by the same virus.

[2] It is interesting to note that in the 1920's two American epidemiologists Reed and Frost demonstrated a discrete *mechanical* model in which coloured balls represented susceptibles and infectives.

by the number of infectives. This basic model assumes random mixing of individuals in the population and does not allow for the complex interactions between physically separated sub-populations, nor for variable incubation or latent periods of a disease. The problems inherent in models that make simplifying assumptions concerning the nature of spatial distributions are discussed in [10]. Various extensions to the SIR model to allow for these phenomena have been added over the last fifty years. Some mathematical models and computer simulations deal with the spatial distribution of susceptibles along a line, across a lattice or over a network to overcome the inaccuracies due to the assumption of random mixing of the population. Cellular-automata and other discretized versions of the SIR method have been utilized also [11, 12]. Some of these models have also incorporated disease *carriers* (e.g. some viruses are transferred by mosquito), and non-homogeneous populations. The model presented in the current paper allows all of these phenomena to emerge from the simulation without hard-coding their behaviour.

The current threats of biological warfare and terrorism have raised the stakes in Western society for epidemiology. The U.S. National Institute of General Medical Sciences has devoted $1.6 billion to a fledgling agent-based study of epidemics [13]. Like the U.S. project, this paper adopts agent modelling to represent the principles of epidemiology in an intuitive but realistic fashion. As shall be shown, the process of epidemic spread offers a means of increasing the genetic and phenotypic diversity of a population and of capping its density.

1.3 Relevant Consequences of Basic Epidemic Theory

There are two theories of epidemiology that are particularly relevant here. The first of these is known as the *Threshold Theorem* [14]: a disease cannot take hold in a population of susceptibles unless the population density is above a particular threshold. This value relates to the infectivity of a disease and the death and recovery rates it induces. If population density passes beyond the threshold, the disease will reduce the population to a level as far below the threshold as it was above it prior to the epidemic.

The *Threshold Theorem* has many consequences, one of which has come to be known as *Herd Immunity* [8, pp. 27-31]. This theory indicates that a calculable number less than the full population needs to be immunized to prevent an epidemic. Unfortunately the theory has been shown to provide inaccurate figures in practice, due to its assumption of random mixing in a population. Nevertheless, it highlights an important aspect of epidemics, namely that the spread of a disease is not dependent on the percentage of a population who are immune, but on the contact between susceptibles and infectives. When a population does not mix uniformly, the supply of susceptibles may be similarly irregular.[3] The model presented in this paper does not assume random

[3] For example, if a socio-economic group is immunized against a disease, and these people do not mix randomly with people from other groups, an epidemic may still occur within the latter groups whilst the former is immunized. I.e. sub-group mixing is *important* in considering the spread of a disease.

mixing of a population, rather the agent interactions are emergent from the simulation.

2 An Agent-Based Simulation of Infectious Disease Epidemics

The present simulation runs in discrete time steps during which a population of agents moves freely about a continuous-space, virtual world. The model was originally devised as a part of a generative, interactive artwork (described elsewhere [15]) that exhibits numerous emergent features typical of Artificial Life simulations. The model's essential features are described below.

2.1 Agent Composition, Behaviour and Evolution

Agents are represented visually as coloured boxes. These have a position and velocity on a continuous toroidal surface. Each agent may wander randomly over the space at a speed inversely proportional to its volume. During each time step of the simulation, agents expend an amount of energy proportional to their volume to move and metabolise. At each simulation time step, energy is gained by an agent from the environment in an amount proportional to its upper surface area as if each box-top was equipped with a solar cell charging a battery. Agents exhausting their energy supply "die" and are removed from the simulation. Agents also age throughout a simulation and are removed if they reach the end of their lifespan.

Agents perceive their neighbours' positions, dimensions and colours within a limited visual range. An agent may accelerate towards (or away from) a neighbour that it finds attractive (or repulsive) as determined by reference to colour and dimension templates it stores. Each agent stores colour and dimension templates marking properties it finds attractive in its partners and templates marking repulsive properties. The closer the match of a particular template the greater the tendency of the agent to seek or flee the neighbour that exhibits it. These tendencies are used to adjust the velocity of the agent as it moves.

If two agents' bodies intersect one another, find one another attractive, and pass a maturity/age threshold test, they may produce a single offspring agent per time step at their current location. The offspring is initiated with energy donated by each of the parents. This donation costs parents an amount of energy specified in their property list. The characteristics of the offspring are specified by the crossover and mutation of the parents' genotypes. This is an array of floating-point values coding the properties listed in Table 1. The system employs a single crossover point and mutation of one gene in every offspring by a random amount between +/- 5%.

New births are subject to an overflow test of the available simulation space. If a birth would cause an overflow the request is refused. Following an unsuccessful request, a random member of the population may be eliminated from the simulation to make room for future requests.

Colour (R,G,B)	*Colour preference*	*Colour abhorrence*
Dimension (X,Y,Z)	*Dimension preference*	*Dimension abhorrence*
Visual range	Offspring energy donation	Lifespan

Table 1. Floating-point agent genotype contents (italics indicate vector quantities).

2.2 Disease Behaviour and Evolution

The agents in the model may carry virtual diseases, transmit them to other agents and succumb to infection themselves. The diseases in the simulation co-evolve alongside the agent population but may only exist *within* a host agent i.e. disease does not persist in the environment. A susceptible agent is exposed to a disease when it intersects with an infective agent. An agent that is carrying a disease cannot be infected by a second disease (i.e. an active disease blocks secondary infection).

If an agent is not carrying a disease, its susceptibility is determined by the match between its own colour and the *colour-signature* of the carried disease to which it is brought into contact. The closer the match between the agent's colour and the colour-signature template of the disease, the higher the probability the disease will infect the susceptible agent during a time step of contact. Simulation diseases also possess a *devastation* value that measures the virulence of a disease. This parameter is used to scale the probability of infection and the amount of energy required of a host to survive a time step of infection.

A parameter determines the lifespan of a simulation disease in each host. Long-lived diseases require a host to invest substantial amounts of energy to overcome infection. If a disease is overcome without the death of the host, the agent acquires immunity to the strain of the disease by adding it to an *immunity list*. Any further contact with this disease will result in an *immune response* that prevents the disease from infecting the agent a second time. If a disease kills its host, or the host dies for any reason, the disease it carries dies also, irrespective of its lifespan.

Each disease has parameters determining its *latent* and *incubation* periods (see section 1.1). A latent disease does not require energy of its host and is not infectious. During the incubation period the agent is infective but does not exhibit symptoms. Agents may visually detect disease symptoms in neighbours, potentially allowing them to steer clear, however this feature was not utilized in the current experiments.

Real diseases such as viruses replicate and mutate within a host much more rapidly than the hosts themselves reproduce, circumventing the host's auto-immune response. Consequently, it is possible for humans to repeatedly catch viruses such as the common cold and flu. To model this, a simulation disease undergoes reproduction during every time step of its lifespan. Disease reproduction is asexual and may result in mutation of the disease parameters: colour-signature; devastation; lifespan; incubation and latent periods. A parameter that sets the frequency of a disease's mutation during reproduction may itself be mutated. Together these parameters allow the diseases to co-evolve with the more slowly evolving agent population.

Diseases are represented in the simulation as coloured shapes rendered within the box bodies of the agents. Fig. 1 illustrates the visualization scheme employed.

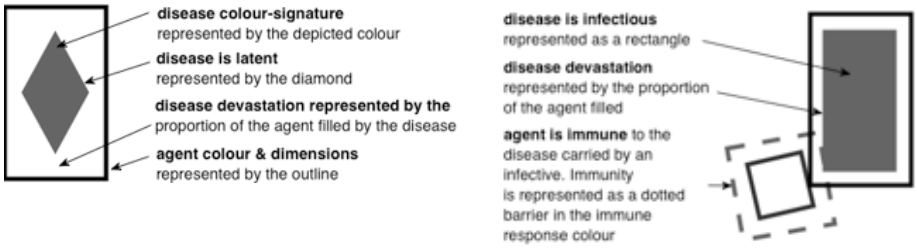

Fig. 1. The visualization scheme for agents and infection.

The parameters for the disease and agents outlined fully specify the features of a epidemic models discussed above. A complex and flexible simulation has been devised that allows for studies of epidemics in non-homogeneous populations with non-random mixing. This agent-based model eliminates many of the problems inherent in earlier epidemiological models.

3 Results

As indicated in the motivation for this work, it had been noted that ecological simulations in which agents competed for resources (including mates and energy) often resulted in a genetically impoverished, homogeneous population. It was hoped that by introducing a novel co-evolutionary model of disease, diversity might be encouraged and uniformity exploited and eliminated by infection.

3.1 Qualitative Discussion

It was found that the disease did indeed exploit the population's uniformity when it arose. Disease also exploited populations of agents that clustered tightly together. In the absence of disease, agents of particular colours and sizes often dominated a simulation, forming large colonies of potential mates. A typical screen shot after 14,000 time steps of the simulation without disease is reproduced in fig. 2(a). Fig. 2(b) illustrates a run after 14000 time steps with identical initial conditions, but in which the disease model was introduced. The diversity in dimensions, colour and spread of the population is far greater in fig. 2(b) than in fig. 2(a). In fact, after as few as 2500 time steps, the non-diseased model often converges to homogeneity and does not break from this condition but drifts gently through genetic space. The population model incorporating disease maintains its diversity indefinitely.

A disease simulation run involves the spontaneous appearance of a disease on average once every one-hundred-thousand agent updates. This new disease is generated with a colour-signature that matches the colour of a randomly selected agent. The agent is infected with the disease and left to continue its travels. Apart from the colour-signature, all other disease parameters for the new infection are randomly generated.

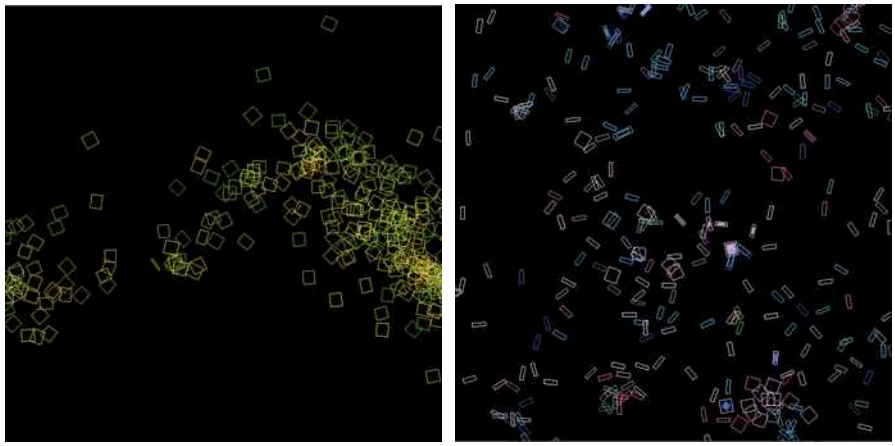

Fig. 2. Two simulation screenshots after 14,000 time steps:
(a) without the epidemiological model; (b) with the epidemiological model.

Depending on the parameters of the new disease, the traits of the infected agent and the population as a whole, the new disease may or may not cause an epidemic. The likelihood of an epidemic is specified by the Threshold and Herd Immunity theories described above. Some observed outcomes are described below along with the conditions giving rise to them in the present simulation environment.

Disease elimination (immediate). If the disease is insufficiently long-lived, or the population is insufficiently dense, or the host does not co-habit with others of a similar colour to itself, then the disease may fail to contact any susceptibles before it dies within the host. The disease will be eliminated from the population immediately.

Disease spread (immediate). A disease may mutate sufficiently within a host to infect susceptibles of a colour significantly different to the original host. If the host mixes amongst others of its kind they may become infected with the disease also. Occasionally the stochastic mechanism allows for a disease to infect a host coloured differently to its own signature. In this case, the devastation of the disease will be low in the infected host but the host nevertheless is able to infect other susceptibles. Such a host may be considered a "carrier" of the disease.

Disease elimination (eventual). If the disease manages to take a hold in the population it may nevertheless die out eventually if the number of susceptibles is reduced. This may happen when a sizeable proportion of the agents encountered by infectives is immune to the disease (even though the population as a whole may not have a significant number of immune members – see footnote 3 above). Circumstances like this arise when agents overcome the disease and acquire immunity, or when the disease is so devastating that it rapidly wipes out the supply of susceptibles before the agents are able to produce many offspring.

Disease spread (continual). A disease well-suited to its environment has sufficient lifespan to ensure it is passed from one susceptible agent to another. Such a disease also needs to be sufficiently devastating that it can be transferred successfully, but not so devastating that it kills off its supply of susceptibles. Diseases that fit these criteria

also have to be sufficiently stable to avoid unwanted mutations that would render them ineffective, but sufficiently mutable so that they can keep infecting an evolving population of hosts. The simulation has given rise to diseases that meet all of these criteria and persist in the population for long periods of time.

Of particular interest are diseases that sustain themselves indefinitely when they are able to utilize susceptibles that are prolific breeders. Such diseases are able to spread through contact between mates who seek one another out (sexually transmitted diseases?) and also by contact between a parent and its newly born. Newly born agents may have traits slightly different to their parents so that occasionally one tends to wander off to seek its own preferred companions, taking the disease to infect others. As long as the disease remains latent for a sufficiently long interval, it will not kill or weaken the agent prior to its immigration to a further enclave.

3.2 Gene Diversity Plot Analysis

Figure 3 gives example plots of the red colour gene value of each agent in the population, versus the simulation time step for (a) a healthy agent population and (b) a population in which disease is present. Each simulation commences with a population of randomly generated agents, and therefore figures 3(a) and 3(b) show red gene values to be widely spread at time step 0.

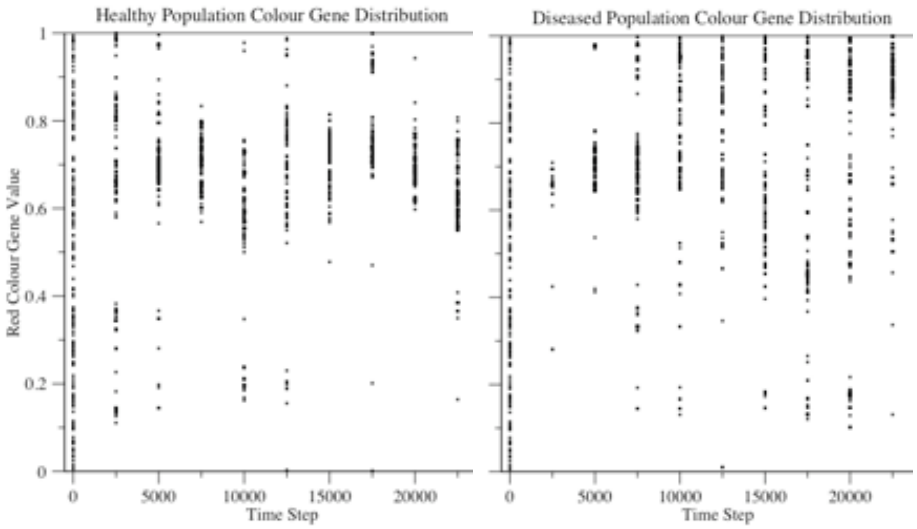

Fig. 3. Plots of all agents' red colour gene values against simulation time step: (a) without the epidemiological model; (b) with the epidemiological model.

Time step 2500 of plot 3(a) commences a long-term decline in the diversity of the red gene in the population. Without the presence of disease, the combination of colour genes an agent possesses determines its mating success based on the presence of potential mates who find the colour of the agent attractive. Thus, agent colour in the disease-free population is driven purely by its ability to attract mates. The decline in diversity is visible as the vertical dispersal of the red colour gene in the population is reduced over time. The few "outliers" at each time step are excursions into new

colours brought out by a momentary success of a sub-population with a specific colouration. Such events may be the result of spontaneous mutations during reproduction. The main population drift in figure 3(a) has red gene values focussed from 0.5 to 0.8.

Figure 3(b) shows the red gene diversity in the diseased population. After the initial random spread of red gene values, the diversity in the population declines dramatically by time step 2500 to a range limited between 0.6 and 0.7. This was the result of a few randomly introduced diseases culling a population that had not yet adapted to existence *without* the presence of disease, i.e. they could not yet locate mates and produce offspring efficiently. A few rapidly acting random diseases therefore wiped out much of the population before it had a chance to evolve strategies for sustaining itself. The system has been programmed to generate several offspring automatically from two randomly selected parents (even if they are not close to one another) in situations like this in order to "jump start" the simulation. This has the drawback of starting a population with a limited gene pool.

The situation depicted in figure 3(b) is especially interesting because with the presence of disease, even this limited gene pool (at time step 2500) does not simply drift about genetic space as did the disease-free simulation when it encountered homogeneity. Instead, as can be seen from subsequent time steps of figure 3(b), the diversity of the population actually expands. The co-evolutionary pressure between the disease and the agents ensures that this situation is maintained indefinitely.

Figure 3 can therefore be seen to confirm the discussion at the beginning of this section and the interpretation of figure 2 given above. In summary, the disease acts to maintain colour diversity in the population, despite pressure applied by mating preferences to the contrary. The disease also forces the population to spread across the available space and it allows the agents to explore a wider variety of shapes than the pressure of the environment alone would have permitted. Both these latter results are clearly depicted in figure 2.

4 Conclusions and Future Work

A model of epidemics has been introduced to an evolutionary, agent-based simulation. The model improved the overall diversity of the population as desired and also encouraged its spread across the available virtual space. A wide variety of disease outcomes emerged from the simulation, each an apparently plausible model of real-world outbreaks.

Future work of interest to the author is a full investigation of the impact of the Threshold Theorem utilizing the present simulation. Can its behaviour be predicted mathematically and demonstrated successfully using this model? It would also be interesting to conduct experiments that model known infectious diseases and their dispersal based upon known interactions of animal or human populations.

References

1. Dorin, A.: The Virtual Ecosystem as Generative Electronic Art. 2nd European Workshop on Evolutionary Music and Art, Applications of Evolutionary Computing: EvoWorkshops 2004. Coimbra, Portugal, Springer-Verlag (2004) 467-476
2. Yaeger, L.: Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision and Behavior or Polyworld: Life in a New Context. Artificial Life III, Addison-Wesley (1992) 263-298
3. Todd, P.M.: Artificial Death. Jahresing 41 (German modern art yearbook) (1994) 90-107
4. Mascaro, S., Korb, K.B., Nicholson, A.E.: Suicide as an Evolutionary Stable Strategy. Advances In Artificial Life, 6th European Conference. Prague, Springer (2001) 120-133
5. Ray, T.S.: An Approach to the Synthesis of Life. Artificial Life II. Santa Fe, New Mexico, Addison Wesley (1990) 371-408
6. Cartlidge, J., Bullock, S.: Caring Versus Sharing: How to Maintain Engagement and Diversity in Coevolving Populations. 7th European Conference on Artificial Life. Dortmund, Germany, Springer-Verlag (2003) 299–308
7. Hillis, W.D.: Co-Evolving Parasites Improve Simulated Evolution as an Optimization Procedure. Artificial Life II. Santa Fe, New Mexico, Addison-Wesley (1990) 313-324
8. Ackerman, E., Elvebeck, L.R., Fox, J.P.: Simulation of Infectious Disease Epidemics. Charles C Thomas Springfield, Illinois (1984)
9. Bailey, N.T.J.: The Mathematical Theory of Infectious Diseases and Its Applications. 2. Charles Griffin & Company London (1975)
10. Durrett, R., Levin, S.: The Importance of Being Discrete (and Spatial). Theoretical Population Biology 46:3 (1994) 363-394
11. Willox, R., Grammaticos, B., Carstea, A.S., Ramani, A.: Epidemic Dynamics : Discrete-Time and Cellular Automaton Models. Physica A 328 (2003) 13-22
12. Martins, M.L., Ceotto, G., Alves, S.G., Bufon, C.C.B., Silva, J.M., Laranjeira, F.F.: A Cellular Automata Model for Citrus Variagated Chlorosis. eprint arXiv : cond-mat/0008203 (2000) 10
13. Pilot Projects for Models of Infectious Disease Agent Study (Midas). Webpage. National Institute of General Medical Sciences (NIGMS) http://grants.nih.gov/grants/guide/rfa-files/RFA-GM-03-008.html (Accessed: October 2004)
14. Kermack, W.O., McKendrick, A.G.: Contributions to the Mathematical Theory of Epidemics. Proceedings of the Royal Society, London 115 (1927) 700-721
15. Dorin, A.: Artificial Life, Death and Epidemics in Evolutionary, Generative Electronic Art. In Rothlauf, F. (ed.) Evoworkshops 2005, Springer-Verlag, Berlin; Heidleberg (2005) 448-457