New styles of thinking for the era of organic networks

M Resnick

When people see patterns in the world, they often assume that there is some type of centralised control, even when it does not exist. This 'centralised mindset' could act as a braking function on the spread of organic networks — if people are not able to think about and understand the decentralised interactions that underlie organic networks, they will not be able to conceive, design, implement, and sustain such networks. In this paper, I discuss how we can help people move beyond the centralised mindset, helping them develop more 'ecological' styles of thinking that are more suited to the era of organic networks.

1. Introduction — the centralised mindset

The papers in this section describe the growing importance of organic networks — decentralised, low-infrastructure networks that rely on local communication rather than top-down centralised control. We can apply the idea of organic networks to many different domains, including communications, management, education, and healthcare. In each of these domains, new technologies provide an opportunity to reorganise activities, relying on decentralised models much more than was previously possible.

But there is an important obstacle to the spread of organic networks — most people are not very good at thinking about and understanding the decentralised interactions that underlie organic networks. At some deep level, people have strong attachments to centralised ways of thinking. When people see patterns in the world, they often assume that there is some type of centralised control, even when it does not exist. For example, most people assume that birds in a flock play a game of follow-the-leader — the bird at the front of the flock leads, and the others follow. But that is not so. In fact, most bird flocks do not have leaders at all. Rather, each bird follows a set of simple rules, reacting to the movements of the birds nearby it. Orderly flock patterns arise from these simple, local interactions. The bird in front is not a 'leader' in any meaningful sense — it just happens to end up there. The flock is organised without an organiser, co-ordinated without a co-ordinator. Yet most people continue to assume the existence of a 'leader bird'.

This assumption of centralised control, a phenomenon I call the centralised mindset [1], is not just a misconception of the scientifically naïve — it seems to affect the thinking of nearly everyone. Until recently, even scientists assumed that bird flocks must have leaders. It is only in recent years that scientists have revised their theories, asserting that bird flocks are leaderless and self-organised. A similar bias toward centralised theories can be seen throughout the history of science.

most people are not very good at thinking about the decentralised interactions

The centralised mindset could act as a braking function on the spread of organic networks. If people do not understand the core ideas underlying organic networks, they will not be able to conceive, design, implement, and sustain such networks. In this paper, I discuss how we can help people move beyond the centralised mindset, helping them develop new ways of thinking more suited to the era of organic networks.

2. The walking tree

To provide a framework for thinking about organic networks, let me start with a simple example. In the rain forest of Costa Rica, there is an unusual type of tree known as a 'walking tree'. It is a strange looking tree. At the base of the tree is a tangle of roots, rising about a metre above the ground. It looks as if someone yanked the tree straight up out of the ground, leaving about a metre of its roots exposed above ground level.

According to rain-forest guides, the walking tree actually changes its location over time (although very slowly). How does the tree move? The roots act as a type of evaluation system, searching for good soil for the tree. If there is good soil on the north side of the tree, the roots on that side dig in deeply and hold firmly. If the soil on the south side is not as good, the roots on that side remain shallow and weak. As the roots on the north side become stronger and deeper, the whole tree gradually shifts towards the north, pulled by the strong roots in that direction. As the tree moves, new roots grow around the new location, some of them extending even further to the north. If the roots find even better soil there, the whole tree will, over time, shift even more to the north. Or, if there is better soil to the east, the tree will slowly shift to the east.

We might say that the walking tree follows a TREE strategy:

- **T**est **R**andomly (send out roots in all directions),
- **E**valuate (determine which roots find the best soil),
- **E**lect (choose which direction to move, based on the information from the roots).

The walking tree executes this strategy over and over; as it moves, it continually sends out new roots to search the area around its new location. Over time, it moves in the direction of better soil. Of course, the walking tree does not actually ‘choose’ or ‘decide’ which way to move, as a person would. But it is useful to think of the tree as executing a type of strategy or algorithm[1].

3. Ecological thinking

The TREE strategy is representative of a broader class of strategies that I call ‘ecological strategies’. These strategies are very common in the biological world, used not only by walking trees but many other plants and animals as well. Ecological strategies share two common characteristics:

- responsive to local conditions — in ecological strategies, decisions (e.g. which direction to grow roots) are based on local information, not centrally planned solutions,
- adaptive to changing conditions — as conditions change (e.g. deterioration of soil on one side of a tree), ecological strategies adjust and produce new solutions tuned to the new conditions; there is no pre-planned script, decisions and solutions change over time.

Ecological strategies might seem inefficient and indirect, but they tend to be simple, flexible, and robust. Many ecological strategies employ decentralised approaches, relying on small contributions by (and interactions among) many simple entities (e.g. the roots of the tree), rather than a single, sophisticated decision-making entity.

[1] The story of the walking tree, which I heard from a rain-forest guide in Costa Rica, is actually a mixture of fact and fiction. The tree (genus Socratea, species .exorrhiza) does move several centimetres during its lifetime, but its movement is not actually guided by soil quality. Nevertheless, the story of the walking tree provides a nice metaphor.

Although ecological strategies are most commonly associated with the biological world, they can be useful in a wide variety of other situations — for example, designing management and organisational structures, solving mathematical problems, co-ordinating communications systems. In the 1940s and 1950s, the field of cybernetics aimed to apply ecological-style strategies to many different types of systems —biological, social, scientific, and technological [2, 3]. The field attracted engineers, biologists, psychologists, and anthropologists, all aiming to forge connections between their disciplines. Cybernetics never developed into a mainstream discipline, but cybernetic ideas have attracted renewed interest during the past 15 years, as part of research efforts in the fields of complex systems [4, 5] and artificial life [6, 7]. In particular, researchers are using TREE-like evolutionary models to describe phenomena in a very broad range of domains [8].

As some scientists have pointed out (e.g. Pagels [9]), these new research initiatives can be viewed, in part, as a reaction against the metaphors of Newtonian physics that have dominated the world of science for the past 300 years. Newton offered an image of the universe as a machine, a clockwork mechanism ruled by linear cause and effect. Today, some researchers are shifting metaphors, viewing their objects of inquiry less as clockwork mechanisms and more as ecosystems[2]. Ideas from ecology, ethology, and evolution are spreading beyond their disciplinary boundaries, influencing other academic fields as well as business practices and public policy [10—13]. The market-place in pollution credits among manufacturers, the Associates program of Amazon, the ratings system in eBay, discussion groups on Slashdot, and the open-source design of Linux software are all based on decentralised, ecological strategies.

ecological strategies require not just a change in lesson plans, but a change in mindset

Despite the growing interest in ecological strategies in science, business, and government, school curricula (particularly at the pre-college level) have been largely unaffected and unchanged. It is rare enough for biology classes to explain or model biological phenomena (such as ant foraging or bird flocking) in terms of ecological strategies; it is far, far rarer for ecological ideas to be used as a basis for design or problem-solving strategies in other (non-biology) classes. Even as educational reform efforts have placed increased emphasis on ‘problem solving’, there has been little emphasis on ecological-style problem solving. Some ‘systems thinking’ approaches (e.g. Senge [14]) incorporate ecological strategies, but few schools have embraced these approaches.

[2] This physics versus biology dichotomy is, of course, overly simplistic. Some areas of physics (most notably statistical mechanics, which focuses on the patterns that arise from local interactions among large numbers of elements) are very related to ideas underlying ecological thinking. And some areas of biology are based on centralised strategies. But the core ideas of ecological thinking (especially those related to adaptation and change) are most deeply rooted in the biological fields of ecology, ethology, and evolution.

Adopting ecological strategies requires not just a change in lesson plans, but a change in the mindsets of students, teachers, and curriculum developers. New approaches are needed to help people become 'ecological thinkers'.

This paper focuses especially on pre-college education, based on the belief that pre-college education offers the greatest point of leverage for changing people's mindsets. But the paper aims to develop foundational ideas that are relevant for learners of all ages. The challenges (and importance) of understanding decentralised systems and ecological strategies are similar for 16-year-olds and 60-year-olds, for students and managers.

In the paper, I probe the nature of ecological thinking, with special emphasis on the use of ecological strategies in design and problem solving (not just as a framework for explaining phenomena in the natural world). Through a set of specific examples, I suggest several different categories of ecological strategies. The goal is to begin to develop a framework that can help people understand the nature and uses of ecological thinking. The paper does not aim to provide a rigorous analysis of the effectiveness and limitations of ecological thinking. Rather, it is more of an essay — intended to provoke thought and to draw attention to styles of thinking that have, too often, been overlooked and undervalued.

Many of the examples in this paper involve the use of computers (and computer networks). That is no accident. Ecological strategies often require repeated application of simple rules and/or massively parallel interactions among many entities. Computers (and computer networks) are particularly well-suited to such tasks. Indeed, it is fair to say that ecological strategies are 'truly computational' [15]. Most uses of computers in math/science education involve a re-implementation of traditional strategies that were previously implemented (albeit less efficiently) with paper and pencil. For example, system dynamics programs like Stella [16, 17] are based on the same differential-equation representations that have long been used by mathematicians and scientists for studying the behaviours of systems. Ecological strategies are different; they involve very different representations and approaches than were traditionally used in the paper-and-pencil era. The point is not that computer-based ecological strategies allow people to do things that they could not do before (though that is certainly true); the point is that such strategies allow people to do things that they never even thought to do before.

most uses of computers in education involve a re-implementation of paper and pencil strategies

This link between computers and ecological thinking is ironic. The popular perception of computers places them in opposition to the natural world, but it is possible that computers could, in fact, lead to a much more widespread application and appreciation of strategies from the natural world.

4. Getting to the roots of the problem

Inspired by the story of the walking tree in the Costa Rican rain forest, I decided to introduce the idea of ecological thinking at a workshop for high-school teachers in Costa Rica. There were about 30 teachers at the workshop, many of them with backgrounds in math and science. Only a few of them had heard of the walking tree, and none of them knew how it moved. After explaining the walking tree's strategy, I suggested that we could use a similar strategy to solve math problems. I explained that each teacher could act like one of the roots of the tree, and collectively they could solve a math problem. The idea was for each teacher to do something quite simple (just as each root of the walking tree does something quite simple), but for the group (like the overall tree) to accomplish a meaningful goal.

I proposed the following algebra problem:

$$2x^2 - 7x + 29 = 3104$$

Typically, students learn to solve such problems by using the quadratic formula:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

for an equation in the standard form of: $ax^2 + bx + c = 0$

Students plug coefficient values into the formula and calculate the solutions for x. In this case, $a = 2$, $b = -7$, and $c = -3075$ (i.e. 29 – 3104).

The quadratic formula certainly yields the correct solutions for this problem. But there is a different, more ecological approach for solving this problem. To illustrate this alternative approach, I asked each of the teachers to pick a random number between 0 and 100. Then, I told them to calculate the left side of the equation, using their randomly chosen number as the value for x. For example, if a teacher chose the number 3, then the calculation would be $(2 \times 3 \times 3) - (7 \times 3) + 29$, for a result of 26. Next, I told them to compare their result with the right-hand side of the equation (3104). Of course, there was very little chance of an exact match: after all, they had chosen their values for x randomly. For the teacher who chose the value of 3, there would be an 'error' of 3078 (i.e. 3104 – 26).

After the teachers had done their calculations, I asked whether any of them had 'errors' of less than 1000. Three of them raised their hands. One had chosen the number 44 for an error of 489. Another had chosen the number 35 for an error of 870. The third had chosen the number 40 for an error of just 155.

I explained that these three teachers represented the 'strong roots' of our tree. They had found good soil, and they should dig in deeply, pulling the tree in their direction. The other teachers had weak roots. They should pull up their roots and pick new numbers close to the strong roots (but not exactly the same). Since the strong roots ranged from 35 to 44, the teachers should all pick numbers between 30 and 50, and repeat the activity.

This time, several teachers chose the number 41 and got an exact match (that is, an error of 0). So the number 41 is a 'very strong root' of the equation $2x^2 - 7x + 29 = 3104$. In effect, the roots of the tree had become the roots of the equation.

Just like the walking tree, the teachers had used a TREE strategy — test randomly (each teacher chooses random numbers as possible values for x), evaluate (each teacher plugs his/her number into the equation and calculates the 'error'), and elect (the group selects the numbers that generated the smallest 'errors'); then repeat the whole process (using numbers close to those elected in the previous round).

For each teacher, the strategy was quite simple to execute. Each teacher performed only simple arithmetic operations (multiplication, addition, and subtraction). But the whole group, working together, was able to solve an algebra problem[3].

Of course, the TREE approach to this problem has some clear disadvantages and limitations. Plugging numbers into the quadratic formula is a more 'efficient' strategy, producing a solution more quickly; it is hardly convenient to gather together 30 colleagues every time you want to solve an algebra problem. Also, the quadratic formula always gives an exact solution; the TREE approach does not.

If I had given the teachers a problem with irrational roots, they never would have converged on an exact solution using the TREE strategy; and if the problem had imaginary roots, the TREE strategy would not work at all. Moreover, the TREE strategy is not fully specified: depending on how the participants 'test randomly' (how do they choose their next number?), the strategy might not converge on an answer.

These are significant limitations. On the other hand, there are several reasons for introducing teachers (and students) to the TREE strategy.

- Usefulness

  The TREE strategy can be applied to wide range of problems. What if I had given the teachers a problem involving a fifth-degree (or other higher-order) polynomial? The quadratic formula would no longer work. Indeed, they could no longer look up a closed-form solution. But the TREE strategy would still be useful. And, as I will discuss later, the TREE strategy (and other ecological strategies) can be useful in many other (non-algebra) situations. The TREE strategy is not just a 'math trick' but a general problem-solving strategy.

- Robustness

  The TREE strategy tends to be more robust than the strategies traditionally taught in schools. Even if one of the participants makes an arithmetic mistake (or, in the case of an actual walking tree, if one of the roots breaks or wanders off in the 'wrong' direction), the overall group would still reach the same result. Small errors generally do not cause big problems in the TREE strategy. The TREE strategy also tends to be robust in another sense: people are likely to forget the details of the quadratic formula, but they are likely to remember the TREE strategy.

- Computerisation

  The TREE strategy is well-suited for computerisation — so it is not really necessary to gather together 30 friends to use the strategy. The Costa Rican teachers, after enacting the TREE strategy themselves, created a computerised version of the strategy written in StarLogo [1, 18], a massively parallel version of the Logo programming language [19]. They created 100 StarLogo turtles[4], and programmed the turtles to follow rules similar to those the teachers had followed in solving the algebra problem. Each turtle started at a random x-position, and it used its x-position as the x-value for the equation. As the turtles executed the TREE strategy, they gradually converged to the same x-value, indicating a solution to the equation. One reason that the TREE strategy has been underused and overlooked in classrooms is that it involves application of the same rules over and over again — a tedious task for humans. Widespread availability of computers makes the TREE strategy much more accessible (both practically and conceptually).

- Problem understanding

  Regardless of whether the TREE strategy is 'better' than other strategies for finding roots of a polynomial, it provides learners an additional way of understanding the idea of finding roots. Minsky [20] and others have observed that you do not really understand an idea until you understand it several different ways. Each way of thinking about something strengthens and deepens each of the other ways of thinking about it. Thus, the TREE strategy, in conjunction with other strategies, provides a more robust understanding.

- Decentralised approach

  Perhaps most important, experiences with the TREE strategy can fundamentally change the ways that people look at the world. As discussed earlier, students usually learn top-down, centralised strategies for solving problems and designing artefacts. The TREE strategy represents a more decentralised approach — lots of separate parts, each with very simple behaviours, work together to produce complex-seeming results.

[3] Note that the given equation has two roots: 41 and –37.5. The way the activity was organised, the teachers found only one of the two roots. But with slight modifications to the activity, allowing negative numbers and non-integers, the teachers could have found both roots, though the process would have taken longer.

[4] Turtles are computational objects with position and heading. With StarLogo, users can write programs for thousands of turtles, then observe the patterns that form from all of the interactions. You can download StarLogo from http://www.media.mit.edu/starlogo/

5. Ecologies on the Net

In recent years, education-research conferences have been full of sessions about the educational implications of the Internet. Many researchers focus on how the Net will provide students and teachers with easy access to huge libraries of information. Other researchers focus on how the Net will make possible new types of learning communities, connecting people with shared interests from all over the world. But little is heard about what, in my mind, is the most important implication of the Internet: how the Net can support and encourage new ways of thinking — in particular, ecological ways of thinking. To the common metaphors of Internet such as library, highway, and market-place [21], we should add the metaphor of Internet as ecosystem.

As is well known, the Internet is based on a decentralised structure. New computers, new users, and new functionality can be added to the Net without any centralised decision-making. The Net makes possible new types of decentralised collaborations, enabling large numbers of people to work together on shared tasks, such as decoding ciphers, creating on-line libraries of virtual LEGO bricks, or organising databases of genetic sequences. In this way, the Internet makes it possible to leverage the small efforts of the many, rather than the large efforts of the few [11, 22, 23]. Parts of the Internet seem to function as artificial ecologies. For example, Best [24] analysed newsgroups in terms of ecological interactions. In this view, ideas on newsgroups compete with one another for the attention of human readers and posters; certain ideas reproduce and flourish (and even spread to other newsgroups), while others die out.

The ecologies of the Internet could be a particularly fertile ground for the development of ecological thinking because they can be designed, manipulated, and analysed much more easily than ‘natural’ ecologies. As Papert [25] has argued, people learn with particular effectiveness when they are actively engaged in the design and construction of personally meaningful artefacts. The Internet enables people to design and play with ‘ecological artefacts’ to a far greater extent than ever before.

Engaging children in ecological-style thinking was one of the motivations behind MOOSE Crossing [26], an on-line community organised by Amy Bruckman (at the time, a graduate student in my research group). MOOSE Crossing is a multi-person, text-based virtual world in which children not only interact with one another but also collaboratively construct the virtual world in which they interact[5]. In MOOSE Crossing, children (mostly between the ages of 9 and 13) create new rooms and objects — and write programs to control the behaviours of those objects (using a scripting language called MOOSE). For example, a ten-year-old girl created a pet penguin that reacts when other people kiss it, hug it, or feed it. The penguin keeps track of how hungry it is, and it reacts differently to six different kinds of food. Another MOOSE Crossing member created a set of potatoes that obey Mendelian genetics; others built a mega-mall with speciality shops. Children can also take on new personas, e.g. a child might decide to become a munchkin and help others in building a replica of Oz. Studies of MOOSE Crossing have documented the ways in which a community can provide strong support for design and construction activities — and, conversely, the ways in which ongoing design and construction activities can support the development of a stronger sense of community [26].

I have been particularly interested in the ways that children on MOOSE Crossing get ideas for new projects. In any design situation, it is often a good strategy to start by looking at previously developed projects, then to consider variations on these model projects. For this reason, it is important for designers to have easy access to good sample projects. Suppliers of computer-based design tools (such as paint programs and programming languages) often include sample projects with their tools, to help users get a sense of what is possible. But there are some important limitations to these pre-packaged sample projects: they reflect the interests and ideas of only the supplier (not necessarily of the full range of interests of the user community) and they are static (changing only if the supplier ships an updated version).

The collection of sample projects on MOOSE Crossing has a very different (and more ecological) feel. The sample projects are created by the MOOSE Crossing members themselves. In fact, since each object created in MOOSE Crossing is fully inspectable and copyable, each object becomes a sample project for everyone else in the community. If a MOOSE Crossing member sees an interesting object, she can ‘look inside’ the object to see the computer code underlying the behaviour — and, perhaps, create a new version of the object with slightly modified code.

The collection of sample projects in MOOSE Crossing continually changes, always reflecting the current interests of the community — without any centralised control. As with other ecological processes, the collection of sample projects automatically adapts and self-adjusts to changing conditions. If a group of MOOSE Crossing members becomes interested in a particular type of project, the collection of relevant sample projects automatically increases. At one point, for example, some MOOSE Crossing members became interested in magic. One member created a magic wand (and wrote a set of programs to accompany it); another created a ‘generic magician’ as a new player class; a third created a spell book. The spell book was full of simple programs that could ‘cast spells’ on other people in the room. Many children in the community made copies of the spell book, and many added new spells (programs) to their personalised spell books. Someone even created a spell that ran all of the other spells in the spell book — a popular program that was soon copied by many others in the community. One child opened a ‘magic store’ where members of the community could get copies of the latest magic-related objects. Eventually, people started to lose interest in magic, and the number of magic-related objects gradually decreased over time — again, automatically adjusting to the current needs and interests of the community.

The interests of the MOOSE Crossing community are always changing, as new members join and as existing members

[5] MOOSE Crossing is an example of a ‘MUD’ [27]. The first MUDs were created to support on-line versions of the game ‘Dungeons and Dragons’; the acronym stands for ‘multi-user dungeons’. MUDs today are used for many different purposes, not just adventure games.

develop new interests. If an 'outsider' tried to supply sample projects to meet these shifting interests, it would be a very difficult challenge. The ecological strategy of letting the participants themselves create the sample projects does a much better job of staying close to the interests and needs of the community. Ecological strategies are especially useful in situations like this, where the problem terrain is constantly changing. In these situations, responsiveness to local conditions and adaptiveness to changing conditions are particularly important.

My hypothesis (yet to be proven) is that children who actively participate in artificial ecologies like MOOSE Crossing will be better prepared to use (and understand) ecological-style strategies in other situations. As the Internet becomes more deeply integrated into everyone's lives, there will be more opportunities to use such strategies. For example, in MediaMOO [28], a networked virtual world intended for media researchers, one participant proposed an ecological strategy for automatically creating new (and useful) paths within the virtual world. The idea was to write a program that kept track of people's movements. If the program noticed that people often went from one particular room to another particular room (moving through several other rooms in between), it could automatically construct 'short cuts' to allow people to jump directly from initial room to the final destination. Several universities tell stories of how an architect, using a similar strategy to decide on the placement of walkways around the university library, surrounded the library with grass and waited a year to see where the paths developed, and then installed the permanent walkways.

This approach allowed the community itself to decide on the placement of the walkways (in an informal, decentralised way). As the Internet allows more people to become 'architects' of multi-user spaces, it will become increasingly important for people to learn how and when to use such ecological strategies.

6. Ecological learning environments

Ecological thinking can apply to education at several different levels. The previous sections focused on the need to help students develop as ecological thinkers —helping them learn how and when to use ecological strategies (especially in the context of new computational media). This section focuses at a different level, discussing how the ideas of ecological thinking can be applied to the design of learning environments themselves. Just as students too often adopt centralised strategies in trying to solve problems, educators too often adopt centralised strategies in designing learning environments.

In the Computer Clubhouse project[6], a network of after-school learning centres for youth from low-income communities, we have explicitly tried to apply ecological ideas to the design of a learning environment [29]. Computer Clubhouses are designed to help members (ages 10—16) learn to express themselves creatively with new technologies, learning to create their own animated stories, robotic constructions, interactive newsletters, video games, musical compositions, and Web sites. With support from Intel Corp, the Clubhouse network has grown (as of May 2004) to 87 sites in 17 countries.

The activity structures at Clubhouses are based on ecological principles. Adult staff and volunteer mentors play very important roles at Clubhouses, but they do not plan activities in a centralised way. There are no 'classes' in Clubhouses, with a teacher lecturing to all of the Clubhouse members. Rather, ideas spread organically within Clubhouses. When Clubhouse members want to start on a new project, they often begin by looking at samples of previous Clubhouse projects (which are kept on display throughout the Clubhouse), then thinking about variations or extensions that they can work on. After that, their projects continue to evolve through various interactions — interactions with other Clubhouse members, interactions with mentors, and interactions with the media and materials on hand at the Clubhouse. Many projects involve groups of Clubhouse members working together, but we do not explicitly organise members into assigned teams, as is often done in classroom-based collaborative activities. Rather, we try to create an environment in which collaborative groups emerge as a natural part of ongoing activities. To a large extent, that has happened. At Clubhouses, projects and project teams are not fixed entities; they grow and evolve over time. A member or mentor might start with one idea, a few others will join for a while, then some others will leave to start working on a related project.

One Clubhouse project, for example, started with two Boston University graduate students who volunteered as mentors. The two students were both enthusiastic about robotics. Initially, they wanted to organise a workshop to teach Clubhouse members about robots. But we encouraged them, instead, to start by building their own robot at the Clubhouse. Our hope was that Clubhouse members would view the graduate students as fellow learners, not as traditional teachers.

For several days, the two graduate students worked on their own; none of the youths seemed particularly interested. But as the project began to take shape, a few youths took notice. One decided to build a new structure to fit on top of the robot, another saw the project as an opportunity to learn about programming. After a month or so, there was a small team of people working on several robots. Some youths were integrally involved, working on the project every day. Others chipped in from time to time, moving in and out of the project team. The process allowed for what Lave and Wenger [30] have called 'legitimate peripheral participation' — different youths were able to contribute to different degrees, at different times. In general, design teams at the Clubhouse form informally, coalescing around common interests. Communities are dynamic and flexible, adapting to the ever-changing needs of the project and interests of the participants —much as the collection of sample projects adapts in MOOSE Crossing.

As in natural ecologies, diversity is important to this process. At the Clubhouse, we have tried to attract a community of

[6] The Computer Clubhouse is a joint project of the Boston Museum of Science and the MIT Media Lab. The first Computer Clubhouse opened in 1993. With support from Intel Corp, the Clubhouse network has grown substantially across 17 countries. For more information, see http://www.computerclubhouse.org

adult mentors with diverse professional and cultural backgrounds. One reason is obvious — a diverse mix of mentors can better match the diverse backgrounds and interests of the Clubhouse membership. But that is only one reason. There is also an evolutionary argument in favour of diversity within a learning environment. New projects at the Clubhouse emerge through a process related to Darwinian variation and selection. The ‘selection’ of new projects works best when there is a rich ‘variation’ in the combinations of youths, mentors, tools, and ideas [31]. As in natural ecologies, diversity at the Clubhouse leads to a greater robustness and adaptiveness in the types of activities and projects that evolve.

as in natural ecologies, diversity is important to this process

Designing an ecological learning environment requires a shift in traditional ways of thinking about ‘control’. Learning experiences cannot be directly controlled or planned in a top-down way. Indeed, the specific experiences at the Clubhouse have been quite different from what we (as developers) expected. Educational designers cannot control exactly what (or when or how) students learn. In some ways, the design of a new learning environment is like the design of a StarLogo simulation. To create a StarLogo simulation, you write rules for individual objects (such as birds), then observe the large-scale patterns (such as flocks) that emerge from the interactions between the objects. You do not program the patterns directly. So too with the design of learning environments. Developers of learning environments cannot ‘program’ learning experiences directly. Rather, they need to foster and support the interests and creativity of the learners, creating fertile environments in which interesting activities and ideas are likely to grow, evolve, and spread.

7. Conclusions — seeing the forest

This paper has argued that people need to develop as ‘ecological thinkers’ if they are to become full participants in the conception, design, and realisation of organic networks. But the examples and metaphors in the paper run the risk of misinterpretation. At an educational-research conference, for example, I described the foraging strategy of an ant colony as an example of an ecological strategy — the colony as a whole accomplishes complex tasks (bringing food back to the nest) and adapts well to changing conditions, even though each individual ant follows very simple rules and reacts only to local stimuli. In the same presentation, I described the uses of ecological thinking in the design of learning environments like the Computer Clubhouse. At the end of my talk, someone in the audience asked whether an ant colony is really a good model for a learning environment. Do we really want to think of students as ants, each following the same simple rules in an almost mindless fashion?

Certainly not. Students in a classroom are, of course, very different from ants in a colony — or roots on a walking tree. But in all of these cases, the behaviours of the overall system arise, often in unexpected ways, from interactions between the parts of the system. The defining characteristic of an ecological strategy is not the simplicity of the component parts, but the nature of the interactions among those parts. Developing as an ‘ecological thinker’ requires a sensitivity to the role of interactions in a system — and an understanding that effective solutions are not always prescribed in a centralised way but rather arise indirectly from many interactions.

Another possible misinterpretation involves the role of non-ecological strategies. This paper highlights the fact that ecological strategies have been overlooked and undervalued in the past. But it would be a mistake to read the paper as a rejection of ‘centralised’ approaches. Indeed, centralised strategies are often more straightforward to implement, and they are more efficient in some situations (for example, when the context is changing slowly or not at all). An exclusive focus on ecological strategies is no better than an exclusive focus on traditional centralised strategies — just as, in economics, an unyielding commitment to market mechanisms can cause just as many problems as an unyielding commitment to centralised planning. If you focus only on TREE-like strategies, you might not see the forest. There is no doubt that the quadratic formula is still a very useful strategy to learn, and teachers still need to play a centralised role in some classroom situations.

An important research challenge for the future is the development of a more systematic framework of how, when, and why ecological strategies are useful. The goal should not be to ignore or replace traditional strategies, but to expand the repertoire of strategies that people have at their disposal — and to help people learn which strategies (or which mixtures of strategies) are best suited to which situations. Indeed, one of the most important benefits of introducing ecological thinking in classrooms is to help students learn that there are, in fact, multiple ways of thinking about problems.

Acknowledgements

Discussions with Brian Silverman, Seymour Papert, Alan Kay, Rick Borovoy, and Uri Wilensky have been helpful in developing these ideas. Alan Kay has also used the phrase ‘ecological thinking’ to refer to a similar set of ideas. Eleonora Badilla organised the Costa Rican teacher workshop where I developed the ‘walking tree’ example. Andy Begel, Eric Klopfer and Brian Silverman have played major roles in the design and implementation of StarLogo. Amy Bruckman co-ordinated the development of MOOSE Crossing; she and Austina De Bonte provided me with useful insights into children's activities on MOOSE Crossing. The section on ecological learning environments grew out of discussions with Natalie Rusk. I would like to thank Andy diSessa and two anonymous reviewers for their insightful comments on an earlier version of this paper. This research has been supported by the LEGO Group, Intel Corp, the National Science Foundation (grants 9358519-RED, CDA-9616444, and CCR-0122419), and the MIT Media Lab’s Digital Life consortium and Center for Bits and Atoms. The Computer Clubhouse is a joint project of the Boston Museum of Science and the Media Lab. Portions of this paper previously appeared in International Journal of Computers for Mathematical Learning, Vol 8, No 1 (2003).

References

1 Resnick M: ‘Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds’, Cambridge, MA, MIT Press (1994).

2 Wiener N: ‘Cybernetics, or Control and Communication in the Animal and the Machine’, New York, John Wiley (1948).

3 von Foerster H, Mead M and Teuber H (Eds): ‘Cybernetics: Circular, Causal, and Feedback Mechanisms in Biological and Social Systems’, New York, Josiah Macy Jr Foundation (1949).

4 Waldrop M: ‘Complexity: the Emerging Science at the Edge of Order and Chaos’, New York, Simon and Schuster (1992).

5 Gell-Mann M: ‘The Quark and the Jaguar: Adventures in the Simple and the Complex’, New York, WH Freeman and Co (1994).

6 Langton C (Ed): ‘Artificial life’, Redwood City, CA, Addison Wesley (1989).

7 Levy S: ‘Artificial Life’, New York, Pantheon Books (1992).

8 Dennett D: ‘Darwin’s Dangerous Idea’, New York, Simon and Schuster (1995).

9 Pagels H: ‘The Dreams of Reason’, New York, Simon and Schuster (1988).

10 Gladwell M: ‘The tipping point: how little things can make a big difference’, Little-Brown (2000).

11 Johnson S: ‘Emergence: the connected lives of ants, brains, cities, and software’, New York, Scribner (2001).

12 Barabási A: ‘Linked: how everything is connected to everything else and what it means’, New York, Perseus Books (2002).

13 Strogatz S: ‘Sync: The Emerging Science of Spontaneous Order’, New York, Hyperion (2003).

14 Senge P: ‘The Fifth Discipline: The Art and Practice of the Learning Organisation’, New York, Doubleday/Currency (1990).

15 Resnick M: ‘Learning through computational modeling’, Computers in the Schools, 14, No 1/2 (1997).

16 Roberts N, Anderson D, Deal R, Garet M and Shaffer W: ‘Introduction to Computer Simulation: a system dynamics modeling approach’, Reading, MA, Addison-Wesley (1983).

17 Doerr H: ‘Stella ten years later: a review of the literature’, International Journal of Computers for Mathematical Learning, 1, No 2, pp 201—224 (1996).

18 Resnick M: ‘Animal Simulations with *Logo: Massive Parallelism for the Masses’, in Meyer J and Wilson S (Eds): ‘From Animals to Animats’, pp 534—539, Cambridge, MA, MIT Press (1991).

19 Papert S: ‘Mindstorms: Children, Computers, and Powerful Ideas’, New York, Basic Books (1980).

20 Minsky M: ‘The Society of Mind’, New York, Simon and Schuster (1987).

21 Stefik M: ‘Internet Dreams: Archetypes, Myths, and Metaphors’, Cambridge, MA, MIT Press (1996).

22 Whitehead S: ‘Auto-FAQ: an experiment in population leveraging’, Proceedings of the Second International WWW Conference, pp 25—38 (1994).

23 Kalil T: ‘Leveraging Cyberspace’, IEEE Communications Magazine, pp 82—86 (July 1996).

24 Best M: ‘An Ecology of the Net: message morphology and evolution in NetNews’, Masters thesis, MIT Media Laboratory, Cambridge, MA (1996).

25 Papert S: ‘The Children’s Machine’, New York, Basic Books (1993).

26 Bruckman A: ‘MOOSE Crossing: construction, community, and learning in a networked virtual world for kids’, Doctoral dissertation, MIT Media Laboratory, Cambridge, MA (1997).

27 Curtis P: ‘Mudding: social phenomena in text-based virtual realities’, Presented at the Directions and Implications of Advanced Computing (DIAC) conference, sponsored by Computer Professionals for Social Responsibility, Berkeley, CA (1993).

28 Bruckman A and Resnick M: ‘The MediaMOO Project: Constructionism and Professional Community’, Convergence, 1, No 1, pp 94—109 (1995).

29 Resnick M and Rusk N: ‘The Computer Clubhouse: preparing for life in a digital world’, IBM Systems Journal, 35, Nos 3 and 4, pp 431—439 (1996).

30 Lave J and Wenger E: ‘Situated Learning: legitimate peripheral participation’, Cambridge University Press, Cambridge, UK (1991).

31 Dawkins R: ‘The Selfish Gene’, New York, Oxford University Press (1976).

Mitchel Resnick explores how new technologies both necessitate and facilitate deep changes in the ways people think and learn. His Lifelong Kindergarten group has developed a variety of educational tools, including the ‘programmable bricks that were the basis for LEGO’s award-winning MindStorms robotics construction kit. He also led the development of StarLogo, a software toolkit for modelling decentralised systems. He is co-founder and principal investigator for the Lab’s Digital Nations consortium. He also co-founded the Computer Clubhouse project, a network of after-school learning centres for youth from underserved communities. He earned a BS in physics from Princeton University, and an MS and PhD in computer science from MIT. Before pursuing his graduate degrees, he worked as a science/technology journalist for Business Week magazine. In 1994, he was awarded a National Science Foundation Young Investigator Award.