Coping with complexity – the role of distributed information in environmental and resource management

Green, D. G.

Environmental and Information Science and Johnstone Centre,
Charles Sturt University, Albury NSW 2640 Australia,
dgreen@csu.edu.au

Abstract

If humankind is to succeed in managing the world's resources, and in conserving the biodiversity of nature, then we need really comprehensive and reliable information - a global information warehouse - from which to plan. Building and maintaining such a system is a huge task. Contributing information needs to be a routine part of the activity of every organisation. Nor can a global warehouse be centralised; it needs to be distributed across every contributing agency. Achieving this goal will require work in three areas. First, data collation requires data standards directed towards crucial problems, the systematic use of metadata, and the coordination of data sources via information networks. Secondly, interpretation will require flexible query systems. These can exploit new object technologies that makes possible distributed processing. Finally, and perhaps the greatest challenge of all, is to interpret vast amounts of information effectively. The new field of complexity theory is beginning to provide relevant principles and methods.

1 Introduction

One of the great challenges facing mankind at the turn of the millennium is how to conserve and manage the world's living and natural resources.

The problem is immense. The planet's surface area totals 509,000,000 square kilometres. Simply monitoring such vast tracts is a huge task. Taxonomists have described about 1.5 million species (Wilson 1992). The total number of species is not known, but is estimated to be somewhere between 10 million and 100 million. At the current pace it would take at least another 300 years of taxonomic research simply to document them all. Modern technology can help with these task (Klomp et al. 1997) but at the same time generates huge volumes of data that must somehow be stored, collated and interpreted.

The problem is also acute. As human population grows the pressure on resources grows with it. We have now reached a point where no virtually no place on earth is untouched by human activity, and where it can be questioned whether the existing resources can sustain such a large mass of people indefinitely. Slowly we are learning to use resources more carefully.

Given the size and urgency of the problem, piecemeal solutions simply will not do. We have to plan and act systematically. And we need sound, comprehensive information from which to plan. The problem is so huge that nothing less than the coordinated efforts of every environment and resource agency in every country will be adequate.

Our ultimate aim should be nothing less than a global information warehouse documenting the world's resources. Until recently such a goal was unattainable. Collating all available information in one place is simply not possible. However improvements in communications, and especially the rise of the Internet as a global communications medium, now make it feasible to build such a system as a distributed network of information sources.

In this account I outline some of the technical issues involved in developing such a system and describe some relevant recent developments.

2 Mobilising data resources

2.1 Information networks

How do we organise information on a large scale? One first approach is to start at the source and organise publishing sites into an information network. In this context an information network is a set of sites on the Internet that coordinate their activities. In particular they operate under some common framework, especially the indexing of the information that they supply.

There are now many networks that focus on environment and resources. As examples some of the networks that I have been involved with include the International Organisation for Plant Information (Burdet 1992), FireNet (Green et al. 1993), and the Australian Biological Research Network (Green and Klomp 1997).

The advantage of information networks is that they can address directly issues that are crucial in building a reliable information system (Green 1995), including:

- standardisation,
- quality assurance,
- indexing, and
- stability.

Information networks operate best where the subject matter can be well-defined and where the operating model can be tightly defined. However the coupling of sites can vary enormously. Databases typically consist of three main components: an interface, data, and tools for search and retrieval. In a distributed database, information and indexes are shared across computers on each of the participating sites (Fig. 1).

A good example of a centralised approach is the popular software resource called TUCOWS. The TUCOWS system is a network of at least 100 software repositories on the World Wide Web. The network has a central reference site and many mirrors, which duplicate varying levels. The incentive for the mirror sites is the "honeypot effect" (Green 1995): the software attracts users.

Perhaps the best examples of well-focussed networks are the Virtual Tourist and the World Wide Web Virtual Library. Both of these services appeared early in the life of the World Wide Web to meet obvious needs. Each system has a central indexing site, but the sources of information about each country or topic

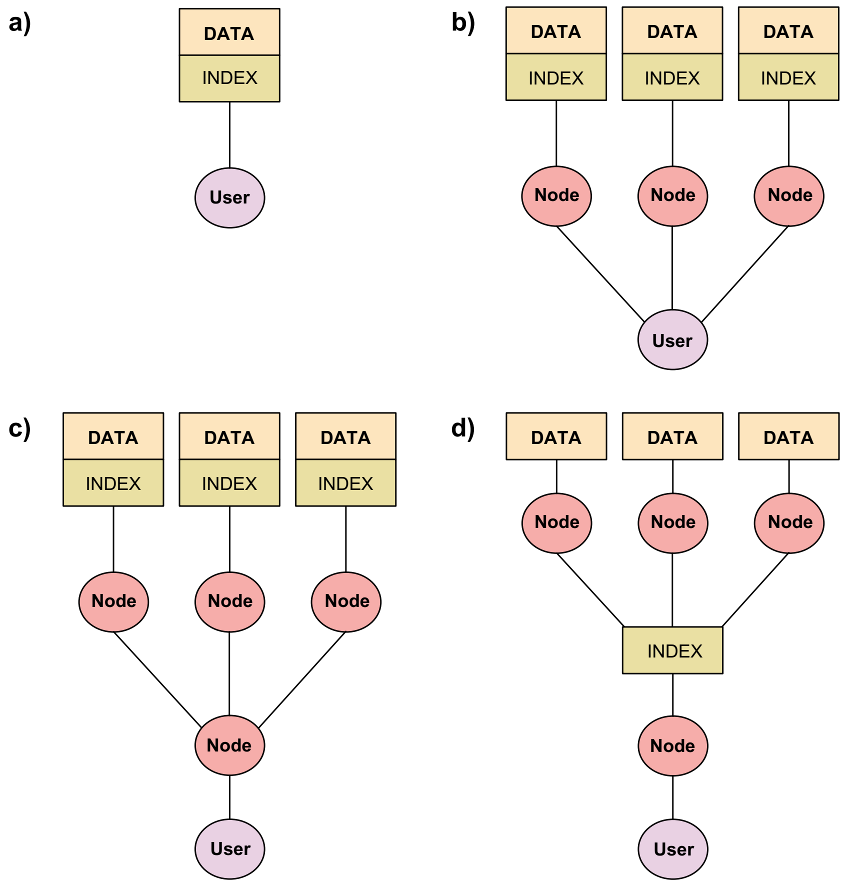

Figure 1. Some alternative designs for distributed data systems. They range from both the index and data being sited on a single machine to various kinds of distributed systems.

are spread across literally thousands of contributing sites.

The current chaotic state of the World Wide Web is a direct result of rapid commercial development. This has put great emphasis on promoting corporate image through Web sites and has discouraged sharing and cooperation between sites. In its efforts to organise on-line material the World Wide Web Consortium (W3C) has concentrated on technical standards that deal essentially with indexing.

2.2 Metadata

Metadata means data about data. Metadata play a crucial role in recording, indexing and coordinating information resources (Cathro 1997).

One approach is to adopt a bottom-up approach by including metadata elements in the head of every on-line document. To this end the Dublin Core defines a set of standard metadata tags to identify the nature, content and provenance of HTML documents. In Table 1, for example are the metadata entries for this document.

The most widespread use of the Dublin Core so far has been in developing metadata indexes for specialised areas. In Australia, for example, the National Library and partners have begun a MetaWeb Project to index government, libraries and some research information. This project also developed software to help automate the addition of metadata to existing documents. There has also been substantial activity towards indexing environmental and resources information. Environment Australia now includes Dublin Core style metadata in all its on-line documents and services. Likewise the National Resource Information Centre now uses DC metadata to index entries for the National Directory of resources information (NDAR), which it has been actively developing for nearly ten years (Shelley 1992).

In order to address the need for metadata for a wider range of material, and to provide greater flexibility, the World Wide Web Consortium has begun developing a more general standard, known as the Resource Description Framework (RDF) (W3C 1998). RDF will use XML (a

```
<META NAME = “DC.creator”       CONTENT = “David G. Green”>
<META NAME = “DC.title”         CONTENT = “Coping with complexity”>
<META NAME = “DC.subject”       CONTENT = “GIS complexity biodiversity “>
<META NAME = “DC.subject”       CONTENT = “environmental informatics”>
<META NAME = “DC.publisher”     CONTENT = “METLA, Finland”>
<META NAME = “DC.language”      CONTENT = “english”>
<META NAME = “DC.coverage”      CONTENT = “world”>
<META NAME = “format”           SCHEME  = “IMT” CONTENT = “text/html”>
<META NAME = “DC.identifier”    SCHEME  = “URL” CONTENT = “http://”>
<META NAME = “DC.date” type=”creation” scheme=”ISO” CONTENT=”1998-04-26">
```

Table 1. Metadata entries for this document.

reduced form of SGML) as its transfer syntax. One advantage of XML's content based tags is that they are effectively self-documenting. The aims of RDF (still in its early stages at the time of writing) are:

- to provide a robust and flexible architecture for supporting metadata on the Internet and WWW; and
- to allow different application communities to define their metadata property sets.

At present the following issues seriously hinder the value of metadata as it is currently used on the World Wide Web.

Misuse

A lot of commercial sites have misused metadata by stuffing metadata headers with spurious terms. For example few people will go looking for a page about widgets, but if you stuff it with hidden fields full of popular terms like games, money, or sex, then the number of hits increases dramatically. Alert to this problem of misuse, many search engines now specifically ignore meta-terms when searching document headers!

Semantics

Although standards such as the Dublin Core define the syntax for metadata elements, they do not define the semantics. That is they do not define the range of each field. For example the subject element really requires a hierarchy of keywords and phrases if we are to group items systematically.

Apathy

Most people simply ignore metadata. It takes time. Few people even know what it is, let alone how to do it. But most of all there is no perceived advantage.

2.3 Databases

Indexing items is just a first step in the business of data gathering. The real issue is what questions do we want to answer with our data? After all the aim of combining data from different sources is to enable us to address questions of substance.

If we wish to mobilise data from individual researchers then there is a need to identify exactly what information we require. As an example take the case of soils data. This can be collected in various ways, ranging from point samples with various physical (e.g. pH, substrate, depth) and site attributes (e.g. Latitude/Longitude, elevation) to areas classified by the dominant types. Some advantages of raw point data are that:

- it can be rescaled;
- it can be recombined with other data more easily;
- it can be used to generate classifications;
- it retains details that are lost in interpreted data (e.g. exact pH value rather than pH class) and may prove important at some time in the future.

However there are also disadvantages. For example, raw data usually has to be processed to derive usable information. Also standards are rarely applied universally, and idiosyncratic sampling methods abounds.

Elsewhere I have argued (e.g. Green, D.G. 1994) for an approach that combines the flexibility with simplicity. That is

- identify raw variables of high priority;
- collect raw data if simple mechanisms exist to retrieve and use it.

For example, field ecologists collect all manner of detailed information about plant and animal communities. Often their data includes fields that are specific to their study alone. Nevertheless archiving data about particular sites would be useful for any future studies of the area concerned. However every study also contains some common data that would be of universal interest. For instance a useful unit of biodiversity data would be a list of species found at a particular site at a particular time, together with appropriate physical data for the site itself.

2.4 The human dimension

All the technical innovations in the world count as nought if people refuse to participate. For instance only a small minority of Web authors and publishers currently include metadata in their publications. There are many reasons why people fail to cooperate. Here are some of them:

- **Ignorance**

  People do not know what needs to be done. They do not know that data repositories exist. They do not know where to find relevant resources. They do not know what metadata is and how to add it. They also do not know why metadata is important.

- **Lack of time or resources**
  Activities like submitting index entries or adding metadata take time and can be awkward. The processes need to fast, simple and automated wherever possible.

- **No benefits**
  People are unlikely to make an effort to share data if they see no benefit in doing so.

- **Fear**
  This includes fear of losing control of your own information, fear of others stealing your results, and fear of the technology.

There are many carrots and sticks to encourage people to cooperate. Biotechnology has reached a stage where it is now virtually compulsory for researchers to submit their raw sequence data to international databases, such as Genbank (Bilofsky and Burks 1988) and European Molecular Biology Laboratory (EMBL) (Cameron 1988) as a prior condition for publication in any leading journal. This practice has proved so fruitful that many biotechnology companies make the bulk of their data freely available too.

Research funding agencies are now thinking beyond research papers as the only outputs of research projects. Tying data sharing to research funding is an extremely powerful way of enforcing cooperation. Along with various colleagues I have advocated expanding compulsory submission of research data to all areas of publicly funded scientific research, especially for projects dealing with environment and resources (Green 1994). However such practices require that three conditions are first met:

- establishment of networks of research data repositories;
- identification of key variables in each field of research;
- flexible standards that recognise common practice without inhibiting innovative research.

Along with the above sticks there are many possible carrots. One is to provide a reward in return for cooperation. For example in running various indexing services, such as virtual libraries, I have found no shortage of people wanting to register their online information because doing so gives their work extra exposure. For instance my Mapmaker program (Steinke et al. 1996) allows users to generate maps at the same time as submitting spatial information. This service is currently being expanded into a public domain spatial query system.

Another way of encouraging participation, and conformance to desirable standards is via accreditation. In this context accreditation is a general term for any mechanism for providing recognition or endorsement of information on another web site. In a sense any hypertext link to material on another site is de facto endorsement of that material. However a stronger form of recognition is to provide a badge or label that can be added to the web page concerned. This approach is particularly effective if the endorsing site belongs to an organisation with substantial credibility, such as a society, major institution or official government agency.

3 Query systems

The value of data lies in what you do with it. If our sources of data are distributed, then so too will be the tools for interpreting those data.

3.1 Distributed information processing

The simplest form of distributed processes are index queries. A number of Internet search engines farm out queries to a set of contributing sites. Several standards, such as the Summary Object Interchange Format (SOIF) exist to support this kind of activity.

For more general processing there is much interest among computer scientists in using the Internet as a large parallel computer. For example, some researchers routinely farm out large problems to a series of workstations. One step beyond this is to systematise the process. The programming language Java is an obvious tool for this kind of activity. It was designed for downloading applications across the Internet and includes features for invoking remote methods. At least one language independent standard - CORBA - has emerged for interface definition in a distributed environment (Veen 1997).

Object oriented representations raise many possibilities for distributed information processing. Objects form flexible, self-contained units that package data with the methods needed to process them. Negotiation between objects raises the potential for combining them into new, more complex information processing. In recent work colleagues and I experimented with our own prototype system of distributed query objects (Moir, unpublished). Seeking to build on existing infrastructure, this system passed queries as HTTP messages. Each HTTPD server in the object network provides a query interpreter and an object store. Objects consists of metadata identifiers, input slots, processes and outputs. Objects pass queries from one to another, each carrying out specialised tasks. In work to date performance issues have proved to be the most serious limiting factor.

To understand the significance of such a system, consider a simple example, such as real-time monitoring of fisheries or wildlife. Data on the different components (e.g. fishing boats, currents, weather) might be drawn from different sources. The analysis might be carried out on one site, then the results sent off elsewhere to overlay the predicted movements on top of up-to-date satellite images. The advantage is that sites can provide standard processing, for many different purposes, of the specialised data that they provide.

3.2 Distributed publishing systems

Perhaps a good example of a simple processing object is the engine of the Mapmaker service (Steinke et al. 1996) that I provide on my server. The Mapmaker itself allows users to generate maps of any area in the world, using data from the Digital Chart of the World. Moreover they can add their own point data to create (say) customised maps of survey sites. However because the engine for the system behaves as a standalone object, we are able to pass queries to it from other applications, such as a spatial query systems for Australian towns and environmental bibliographies, and a geographic virtual library.

The above applications are byproducts of a more general project on Web publishing languages (Green 1996, Green et al. 1998). The resulting programming language SLEEP (Newth and Green in prep.) simplifies the installation and operation of many kinds of on-line publications. Interpreting a form is a single function call, for example. We are using prototype versions of the language, which makes extensive use of metadata and SGML/XML, to automate many services on the Johnstone Centre Web site (http://life.csu.edu.au/). Besides automating standard Web operations, the language also makes provision for distributed publications. We have also developed specialised modules for several environmental tasks, including map drawing, spatial queries, data plotting and simple spatial statistics.

3.2.1 Example - great circle distances

A simple demonstration of using publishing languages to develop new applications on-line is the form located at

http://life.csu.edu.au/cgi-bin/itc332/calc?circle

This example determines the great circle distance between any two points that the user selects on a map of the world. The elements needed to drive the example are:

- a map of the world,
- a form,
- the language interpreter,
- the following script:

```
source circle.sgl
    target STDOUT
    var slat  -34.5
    var slong 147
    var tlat   -33
    var tlong 149.2
    radius    6306
    circle
    sub
```

Later versions of the interpreter used in this example (Newth & Green in prep.) have been combined with CSU's Mapmaker program (Steinke et al. 1996) to provide a toolkit for implementing a wide range of on-line spatial queries and functions.

3.3 On-line geographic information systems

The Web is fast becoming a standard platform for geographic information systems (GIS). Commercial developers are already producing on-line versions of their GIS software. Also there are many on-line services that include spatial queries or map drawing. Within Australia alone the range of environmental information available on-line is impressive. The following are simply examples. They include queries, map-making, and spatial metadata:

- Species mapper (Environment Australia) http://kaos.erin.gov.au/search/mapper.html
- Data Locator (Australian National Geoscience Information System) http://www.agso.gov.au/ngis/locator.html
- CSIRO AGCRC GRASSLINKS http://www.ned.dem.csiro.au/AGCRC/4dgm/grasslinks/
- Spatial Query and Retrieve Information (Environment Australia) http://kaos.erin.gov.au/database/db.html
- Australian Spatial Data Directory (ASDD) http://www.nric.gov.au/nric/data/data.html
- Australia New Zealand Land Information Council http://www.anzlic.org.au/anz_site.htm

Interactive GIS on-line really requires Java based functions to allow query answer sessions to proceed at an acceptable pace (Steinke et al. 1996.). Spatial metadata concerns records that can be identified with features or locations in the landscape. Such information is particularly relevant when considering natural resources (ANZLIC).

4 The problem of complexity

It is important to realise that simply setting up an information system will not automatically resolve every question. Assuming that we can overcome all the other difficulties perhaps the greatest challenge still remains. And that is the sheer complexity of the world's environment (Bossomaier & Green 1998, Green 1997). Here we encounter some really sobering truths. In this section I shall briefly explain some of these problems and suggest how they might be addressed.

4.1 The serendipity effect

A standard operation in GIS is the overlay. That is you lay one data layer (e.g. roads) over another data layer (e.g. topography). An important application of this process is to examine the interaction between two distinct systems within a landscape. In effect this procedure asks a simple question: how do the two systems represented by the layers interact with one another.

Now a really comprehensive spatial database might contain dozens of layers. For each layer that we add, we can overlay it on any combination of the existing layers. So the number of possible overlays - in effect the number of questions that we can address with the data - increases exponentially as we add data. Thus the chances of making an interesting discovery by chance – by serendipity – increases exponentially as we add new datasets.

The converse of the above is that the size of any problem that we may need to solve goes up exponentially too. This includes the likelihood of having incompatible data (e.g. different scale, fields etc) and the problem of combinatorial explosions in running classifications and other algorithms.

4.2 Unpredictability

Perhaps the least understood aspects of ecosystems are the ways in which local interactions affect the global composition and dynamics of whole communities. Issues such as spatial heterogeneity, unpredictability and networks of species interactions have a great bearing on the viability of strategies to manage ecosystems, especially in altered landscapes. In trying to address environmental complexity, ecology can benefit from greater dialogue with computing and information science.

New areas of research, including artificial life, evolutionary computation and complexity theory, offer the prospect of answering such questions. These new fields span many existing disciplines and are developing a body of theory and methods for dealing with a wide range of ecological problems. The notion of complexity had its origins in physics but now resides at the interface between biology and computing. This chapter presents a brief overview of the science of complexity, the ways in which it applies to the living world, and some of the new approaches that are emerging to deal with them in practice.

Environmental complexity arises from several sources (Green 1997). First there is the sheer number of organisms and environmental features. But more important are the interactions and combinations between all of these, which are truly vast in number and variety. Finally the interactions are usually non-linear.

Environmental complexity manifests itself in many ways. Some examples include the following.

- Non-linearity leads to two important properties. First there is sensitivity to initial conditions. Small, local variations can lead to major differences in global structure and behaviour.
- Instead of averaging out, as in linear systems, local irregularities or interactions in non-linear systems can expand to affect global properties and behaviour.
- Environments and ecosystems exhibit criticality. Many processes will not occur unless some control parameter reaches a critical value. Examples include the spread of fires, epidemics, and the spread of exotic plants. These processes depend, respectively, on the presence of critical levels of fuel, susceptible animals and vacant growing sites.
- Perhaps the most sobering property is that systems can be inherently unpredictable. That is, no matter how much information we collect, we will never be able to predict the future structure or behaviour of the system.

4.3 Connectivity

An important example of complexity, especially in landscapes, is the critical transition between connected and fragmented distributions. For instance if we remove small patches forest from a landscape then the forest as a whole retains its integrity.

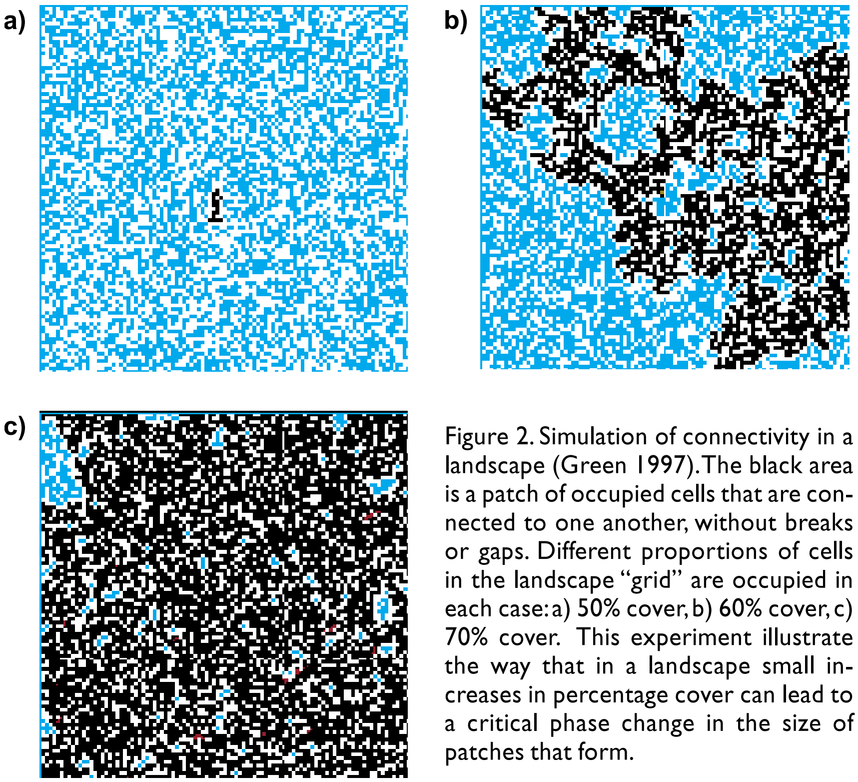

Figure 2. Simulation of connectivity in a landscape (Green 1997).The black area is a patch of occupied cells that are connected to one another, without breaks or gaps. Different proportions of cells in the landscape "grid" are occupied in each case: a) 50% cover, b) 60% cover, c) 70% cover. This experiment illustrate the way that in a landscape small increases in percentage cover can lead to a critical phase change in the size of patches that form.

However if clearing continues (at random), then instead of small patches breaking off the entire system remains connected until a critical point, whereupon it breaks down into many isolated fragments (Fig. 2).

4.4 Modelling and analysis

Although environmental complexity prevents prediction in many circumstances, we can still deal with issues by examining scenarios. For example although we cannot reliably predict the exact nature of global climate change, we can nevertheless examine the potential impact of particular scenarios. Here the advantage of good environmental data becomes clear. For example Environment Australia's Species Mapper not only allows the user to plot maps of species records, but also to generate BIOCLIM and GARP models of their distributions. These models can be run against relevant climate change scenarios to predict the potential effects on plant and animal distributions. As another example, we can run a model of starfish outbreaks on top of a classified satellite image of a coral reef (Fig. 3).

Several systems are now available for modelling complexity in environmental systems. Perhaps the most notable is SWARM (Langton 1998) which provides tools for generating cellular automata and multiagent models.

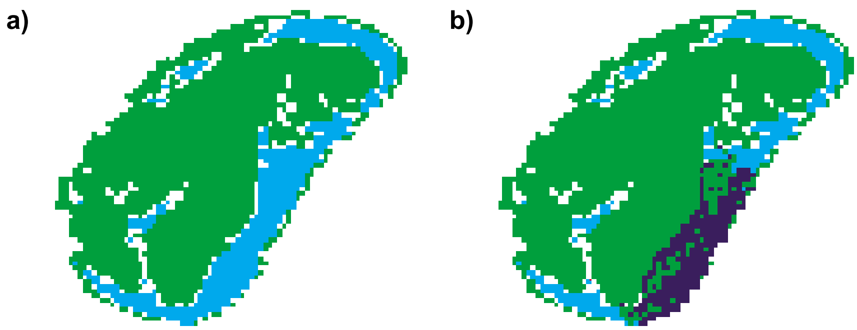

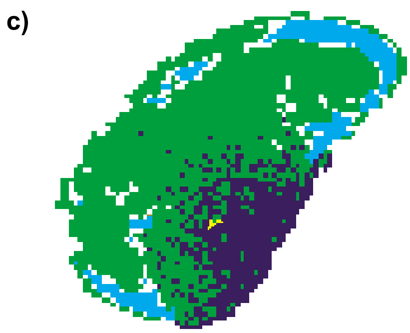

Figure 3. Testing the effect of different scenarios about behaviour of the Crown of Thorns (COT) starfish (Green 1990). In this cellular automaton model, each cell represents an area of reef (or lagoon or ocean) and the dark areas show the sites occupied by starfish. The images show: a) The classified satellite image. b) COTs seek areas of high coral cover. c) The starfish movement is a random walk.

Multi-agent models (Green 1997) teach us that on-line, intelligent agents offer many interesting prospects for handling information (Bossomaier & Green 1998). The most important is the potential for self-organization. Agent-based computing reflects links between computing and biology that are growing ever closer. The organization of an ant colony, for instance, arises out of many simple interactions of ants with each other and with their environment. Likewise software agents on the Web offer the potential to organize information on the Web. For instance an agent acting for an individual user might negotiate with agents at each site that user visits to identify other useful sites and items of information. This raises the potential for "like-minded" agents to form mutual communities, amongst other things.

5 Discussion

As outlined here the paradigm that seems to be emerging is one in which environmental planning, whether local or regional, relies heavily on the interpretation of up-to-date information. Gathering that information will be part of the normal working practice of every agency and every individual researcher. Publishing and indexing that information will be the role of a coordinated network of on-line information servers and Internet publishing sites.

In planning for the future it is important to have a clear image of where you want to get to. What might an environmental information warehouse look like to a user (say) 50 years from now? My vision of the future of the Internet is what I call the "Knowledge Web". That is the present emphasis on sites, home pages and proprietary concerns disappears. Instead the user simply looks for information about a topic and is guided to that information by an intelligent system that actually teaches you as you go. This view would also apply to the proposed environmental information warehouse. Suppose for example that a student wanted to know about conservation of plants and animals in the local area. Starting from some general heading (say plants) the system might guide the student through relevant topics (e.g. biodiversity, conservation, geographic information) at each stage providing background information and links to other information. A geographic query might involve selecting an area on a map and choosing what to be shown from a range of choice offered.

Suppose alternatively that a public servant wanted to see a report about (say) natural resources in southwest Tasmania. After selecting the exact area and topic she might use a report generator to select the kinds of items she wanted to include. The choices might cover a standard list of resources, time period, geographic area, type of material (e.g. policy papers, scientific studies, educational material) and types of items (e.g. maps, tables, graphs, text etc). The system would then build a preliminary report with the option of going back and exploring any aspect in more depth.

The above may seem like science fiction, but it is perfectly feasible. How would it work? Many details would rely on future research and not a few developments in intelligent information systems. However some of the background is clear from our earlier discussion. Everyone would contribute information. Field workers would submit raw data; publishing sites would index and massage and archive it. Information networks would merge it into seamless databases that are organised by topic, location and other relevant keys. They would also provide the query functions, models and other facilities that form the user interface. The interface itself might be built around sets of standard queries, with higher level functions (e.g. the report generator) built from flexible scripts that can adapt to a user's profile and interests.

It is clear that some kind of global environmental information system will need to emerge early in the new millennium. The challenge before us is to ensure that such a system is developed in a consistent, orderly fashion. Whatever its final design and organisation may be, this environmental information warehouse represents the working core of a new paradigm for environmental planning and management.

References

Australian Biological Research Network. http://life.csu.edu.au/abren/

Australia-New Zealand Land Information Council (ANZLIC). http://www.anzlic.org.au/anz_site.htm

Bilofsky, H. S. & Burks, C. 1988. The GenBank genetic sequence data bank. Nucl. Acids Res. 16: 1861-1863.

Bossomaier, T.J. & Green, D.G. 1998. Patterns in the Sand - Computers, Complexity and Life. Allen and Unwin.

Burdet, H.M. 1992. What is IOPI? Taxon 41: 390-392. International Organisation for Plant Information (IOPI). http://life.csu.edu.au/iopi/

Cameron, G.N. 1988. The EMBL data library. Nucl. Acids Res. 16: 1865-1867.

Cathro, W. 1997. Metadata: An Overview. A paper given at the Standards Australia Seminar, August 1997. http://www.nla.gov.au/nla/staffpaper/cathro3.html

CORBA. http://www.corba.org/

Dublin Core. http://purl.org/metadata/dublin_core

ERIN Species Mapper. http://kaos.erin.gov.au/search/mapper.html

Extensible Markup Language (XML). http://www.w3c.org/XML/

Green, D.G. 1990. Cellular automata models of Crown-of-Thorns outbreaks. In: Bradbury, R.H. (ed.). Acanthaster and the Coral Reef - a Theoretical Perspective. Lecture Notes in Biomathematics Series, Springer. p. 157-168.

— 1994. Databasing diversity - a distributed, public-domain approach. Taxon 43: 51-62. http://life.csu.edu.au/~dgreen/papers/taxon.html

— 1995. From honeypots to a web of SIN - building the world-wide information system. In: Tsang, P., Weckert, J., Harris, J. and Tse, S. (eds.). Proceedings of AUUG'95 and Asia-Pacific World Wide Web '95 Conference. Charles Sturt University, Wagga Wagga, p. 11-18.

— 1996. A general model for online publishing. In: Bossomaier, T. and Chubb, L. (eds.). Proceedings of AUUG'96 and Asia-Pacific World

Wide Web '96 Conference, Australian Unix Users Group, Sydney. p. 152-158.

— 1997. Complexity in ecological systems. In: Klomp, N. & Lunt, I. (eds.). Frontiers of Ecology - Making the Links. Elsevier, London. p. 221-231. http://life.csu.edu.au/esa/esa97/papers/green/green.htm

— , Gill, A.M. & Trevitt, A.C.F. 1993. FireNet – an international network for landscape fire information. Wildfire – Quarterly Bulletin of the International Association of Wildland Fire 2(4): 22-30. http://www.csu.edu.au/firenet/fire.html

— & Klomp, N.I. 1997. A Biological Research Network for Australia. Australian Biologist 10(2): 117-120.

— , Bristow, P., Ash, J., Benton, L., Milliken, P.& Newth, D. 1998. Network Publishing Languages. In: Ashman, H. & Thistlethwaite, P. (eds.). Proceedings of the Seventh International World Wide Web Conference. Elsevier, Amsterdam. http://life.csu.edu.au/~dgreen/papers/www7.html

Klomp, N.I., Green, D.G. & Fry, G. 1997. The role of technology in ecology. In: Klomp, N. & Lunt, I. (eds.). Frontiers of Ecology – Making the Links. Elsevier, London. p. 299-309. http://life.csu.edu.au/esa/esa97/papers/klomp/klomp.htm

Langton, C. 1998. http://www.santafe.edu/projects/swarm/

Mapmaker. http://137.166.132.18/cgi-bin/gis/Map/

MetaWeb Project. http://www.dstc.edu.au/RDU/MetaWeb/.

Netscape Communications Corporation 1997. Netscape Publishing System White Paper. http://home1.netscape.com/comprod/products/iapps/capps/pubsys_white_paper.html

Reichelt, R.E., Bainbridge, S. & Green, D.G. 1988. Crown-of-Thorns dispersal in the Great Barrier Reef – a simulation study. Mathematics and Computers in Simulation 30: 145-150.

SGML. http://www.sil.org/sgml/archsite.html

Shelley, E.P. 1992 . Access to data through FINDAR. Proceedings of Workshop: Geographic information systems, cartographic and geoscience data standards, Canberra, 18–20 March 1992. Bureau of Mineral Resources, Australia, Record 1992/27. p. 235-44.

SLEEP. http://life.csu.edu.au/languages/

Steinke, A., Green, D.G. & Peters, D. 1996. On-line environmental and geographic information systems. In: Saarenma, H. & Kempf, A. (eds.). Internet Applications and Electronic Information Resources in Forestry and Environmental Sciences. EFI Proceedings No. 10. European Forestry Institute, Joensuu (Finland). http://sauna.efi.joensuu.fi/~i_dgreen/sql/

TUCOWS. http://www.texoma.com/mirror/tucows/mirror.html

Veen, J. 1997. Object-oriented publishing on the Web. Wired Magazine. http://www.hotwired.com/webmonkey/97/47/index0a.html

W3C 1998. Metadata Home Page: http://www.w3.org/Metadata

— 1998. Resource Description Framework (RDF). Home Page: http://www.w3.org/Metadata/RDF RDF Model and Syntax Working Group: http://www.w3.org/Metadata/RDF/Group

Wilson, E.O. 1992. The Diversity of Life. Penguin, New York.

World Wide Web Consortium (W3C). http://www.w3c.org/

The World Wide Web Virtual Library. http://vlib.stanford.edu/Overview.html