An Object-Oriented Data-Driven Migration Model

Keith Downing
IKU Petroleum Research
Trondheim, Norway
keith.downing@iku.sintef.no

This research uses object-oriented simulation to derive the migratory patterns of animal populations. These patterns are then compared with the trajectories of aquatic pollutant releases to yield quantitative biological impact assessments. These estimates are more realistic than those based on assumptions of uniform or random distributions of animals in a region - a common approach to this problem. The general migration model (MIGMOD) employs stochastic movement routines driven by biological field data to derive heterogenous migration patterns similar to those observed in nature. This data-driven approach circumvents many of the difficulties involved in migratory modelling from first principles. However, we also discuss the importance of selected causal mechanisms in these models.

I. Introduction

To accurately assess the biological impacts of marine oil and chemical spills, realistic models of the migratory movements of indigenous species are needed. Unfortunately, macroscopic migratory patterns are difficult to generate using only macroscopic constructs, since lumped parameter systems abstract away too many important heterogeneities in population distributions and individual behaviors (Hogeweg & Hesper, 1990). Animals have daily and seasonal periods of migration or movement, such as the diurnal trips to feeding grounds of nesting birds, or the spawning runs of fish. These peak movement periods involve large deviations from random or uniform distributions, since the paths taken by migrating animals, particularly those that flock or school, often constitute a very small portion of the total area of the ecosystem. Similarly, sites for feeding, nesting and other activities are often concentrated in small sub-regions of a habitat. Furthermore, assumptions of average migratory patterns of entire populations do little justice to the behavioral contrasts between age/sex classes. This level of detail is very difficult to capture in standard population dynamic models, and is quite problematic for finite-element approaches, since state variables such as local population densities that are modified by transport processes fail to recreate the point to point movements of individuals or small groups.

Since movement patterns and relevant behaviors vary significantly across both species and age/sex classes (Mehlum, 1992), the task of migration modelling entails a commitment to capturing these idiosyncrasies (French & Reed, 1989).

However, this detailed modelling incurs a significant programming burden, since movement routines are required for the many diverse species and behavioral classes. For example, birds typically fly from one site to another and thereby avoid many aquatic pollutant spills, whereas swimming seals may be more vulnerable. Similarly, reproductive female polar bears are less susceptible than young adult males to winter oil spills, since the former den with their cubs on solid ice until spring, while the latter feed along the ice edge.

The modelling/programming dilemma becomes one of managing the trade-off between generality and specificity such that important behaviors are modelled, but flexibility is preserved. The object-oriented paradigm, embodied in programming languages such C++ and Smalltalk, elegantly handles this trade-off. By implementing MIGMOD in C++, we cover a wide range of the generality-specificity spectrum. Hence, MIGMOD can simulate (a) idiosyncratic behaviours of particular species and age/sex classes of animals such as immature female ringed seals, or (b) "generic" animals in cases where extensive age/sex (or even species-specific) information is lacking, or (c) animals from the middle of the general-specific spectrum, such as mature polar bears of unspecified sex. Furthermore, MIGMOD easily incorporates new species and age/sex classes and their associated behaviors.

Modelling of migratory behaviour based on causal processes and environmental stimuli (e.g. Balchen, 1994; Reed and Balchen, 1982) remains too complex for most applications. However, modelling based on observed animal distribution patterns and movements has been successfully applied to marine mammals (Reed et al., 1988) and birds (Ford et. al., 1982). This data-driven approach underlies MIGMOD, which uses field data as a framework for stochastically determining the movements of simulated animals within and between critical habitats. In essence, MIGMOD employs migratory data concerning "where" and "when" to circumvent the problems of modelling "how" or "why".

MIGMOD's data-driven approach is facilitated by the existence of accumulated field data concerning:

- traditional site locations for activities such as nesting, feeding and wintering,
- ranges of arrival times and durations of visits at these sites, and
- general and age/sex-specific migratory behaviors of the animals,

MIGMOD uses this biological field data to predict daily and seasonal migratory patterns for birds, aquatic mammals and polar bears. These patterns are then compared with the projected trajectories of aquatic pollutants to yield an impact assessment that avoids many of the inaccuracies of more abstract models.

II. Primary Data Structures

The migratory patterns generated by MIGMOD result from the interactions of 4 key components: the animal class hierarchy, the habitat grid, site-habitat mappings, and the travel itinerary. The former partitions simulated animals (a.k.a. animal "particles") into age/sex groups that range from the very general *animal* class to specific types such as *reproductive male seal*. The habitat grid partitions the environment into a potentially wide variety of niches, while site-habitat mappings determine the site types that correspond to these niches. Finally, the travel itinerary defines the rules for inter and intra-site movements. Both the site-habitat mappings and the major components of the travel itinerary are animal-class specific, while all animal types share the same habitat grid. Only the animal class hierarchy is predefined; the other 3 components are entered by the user.

II.1. Animal class hierarchy

Figure 1 portrays MIGMOD's animal class hierarchy. Note that the direct descendants of the general animal class are of three types:

1. developmental stages of general animals (dependent, independent),
2. sexual statuses of a general animals (male, female),
3. specific animal types (bear, seal, bird).

The complete animal hierarchy consists of over 50 subclasses, many of which arise via multiple inheritance from the 7 second-level categories. Although specific simulation scenarios usually involve 6 or fewer classes, all levels of the hierarchy are completely simulatable, with more generic levels typically having simpler movement routines.

A typical MIGMOD simulation involves a few hundred individuals from each of several classes. With the exception of parents and their immature children, individual animal objects have no knowledge of the behaviors of other individuals in their own or other classes. As discussed below, each animal object's behavior is determined by the site-habitat mappings and site constraints defined for its particular class.

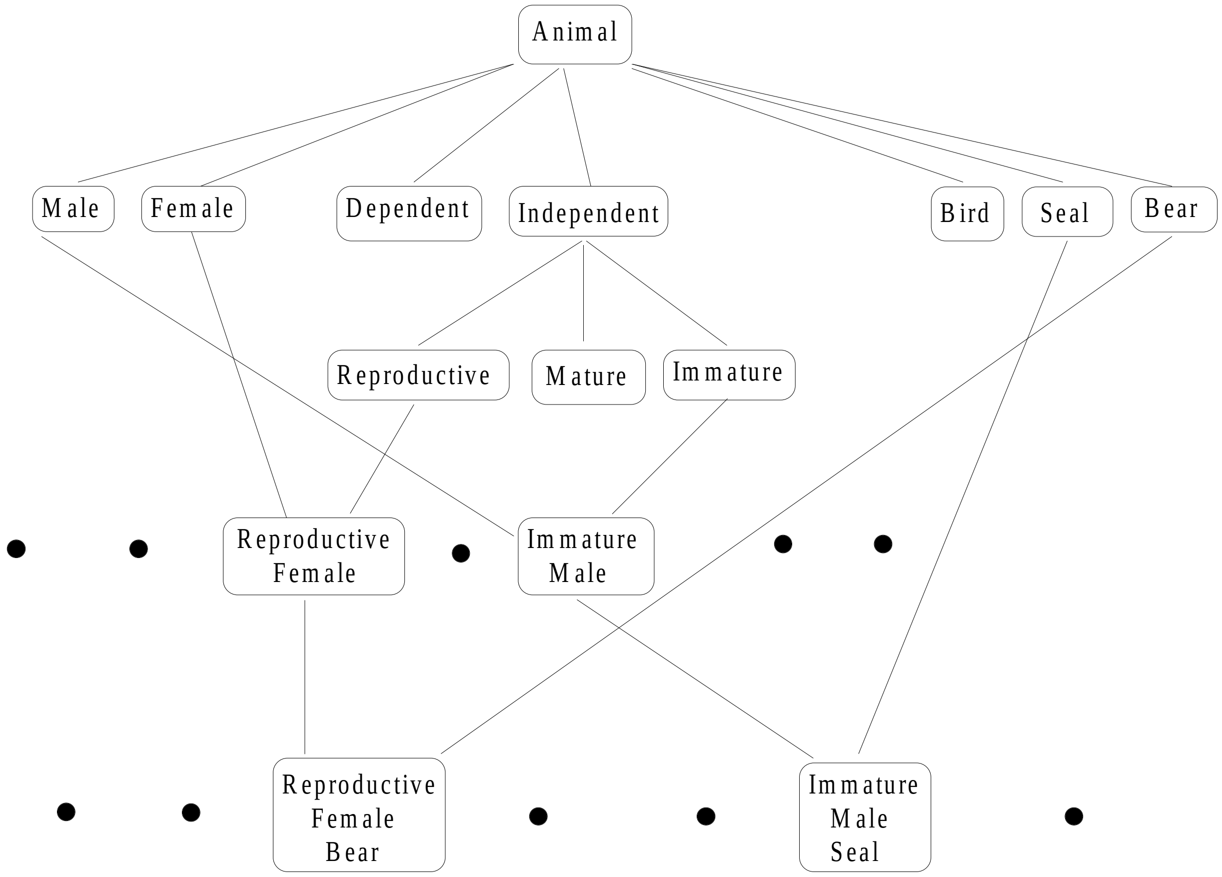

Figure 1: Selected components of MIGMOD's animal-class hierarchy

To review a bit of object-oriented programming, there are two key elements of a parental subclass that are inherited by its child subclasses: property variables and methods. The generic *animal* class contains property variables such as position (i.e., longitude and latitude), migratory speed, an alarm (indicating the departure time from the animal's current site), a travel itinerary (see below), etc. More specific classes such as *reproductives* have additional properties such as a birthing time range and a list of children.

In OOP, "methods" denote subroutines which only apply to objects of the given class and any of its descendant classes. However, the descendant classes have the ability to rewrite/override any of the ancestral methods with a local version of the same name. For instance, all subclasses inherit the *makemove* method from the *animal*, but the *reproductive* subclass has its own version of *makemove* (see below).

II.2 Habitat grids

The simulation environment is divided into a square grid, typically 50 x 50, where each cell represents a particular habitat type. Standard types include "seaward cobble beach", "landward rocky shore", "subtidal sand bottom", etc. All animal movements occur within this gridded area, with the exception of migrations to and from a possible "out of grid" habitat.

II.3 Site-habitat mappings

In MIGMOD terminology, a *site* is a functional classification of a grid cell determined by the activity that a particular animal performs there. Typical site types include feeding, wintering, molting, and breeding. Basically, each animal class employs the same habitat grid but with a different interpretation of its grid cells, as defined by the site-habitat mappings. For example, male birds of a particular species may have feeding sites at *open sea* habitats, while females may prefer *seaward sandy shore* habitats for the same activity. Each simulated animal class has a set of user-defined relevant site types plus a mapping from each site to one or more habitat types. This mapping defines the geographic basis for the movement routines by restricting animals to certain habitat cells during various daily and seasonal behavioral periods.

In conjunction with the site-habitat mappings, the user can define a *tenacity grid* for a particular site type. This gives weights to the cells of the habitat grid to establish a visiting priority among all of the habitat cells that correspond to the site type. For example, if female birds have nesting sites on the *landward rocky shore* habitats of an archipelago, and one island is twice as popular as all others, then tenacity weights would be assigned to all *landward rocky shore* cells such that the favored island's cell's weights would be double the weights of other habitat cells of the same type. These weights will then bias the site selection process. Tenacity grids are typically based on population density distributions derived from field data.

II. 4 Travel itinerary

Each animal particle has a travel itinerary consisting of a *site graph*, *site constraints* and a *travel history*; the former two are shared by all members of the age/sex class, while the latter is particle specific. In the site graph, nodes represent site types, and arcs denote legal migratory moves. For example, in the site graph of Figure 2, migrations from wintering to molting sites are legal, as are forays between different feeding sites (i.e., different habitat cells that correspond to feeding sites under the animal class' site-habitat mapping), but migrations from feeding sites to molting sites are prohibited.

The site constraints (drawn near their associated site nodes in Figure 2) are temporal restrictions upon:

1. The visit season - time range for arrivals at sites of this type.
2. The visit length - minimum and maximum durations of individual visits.
3. The visit frequency - expected number of arrivals (per individual) during the visit season.

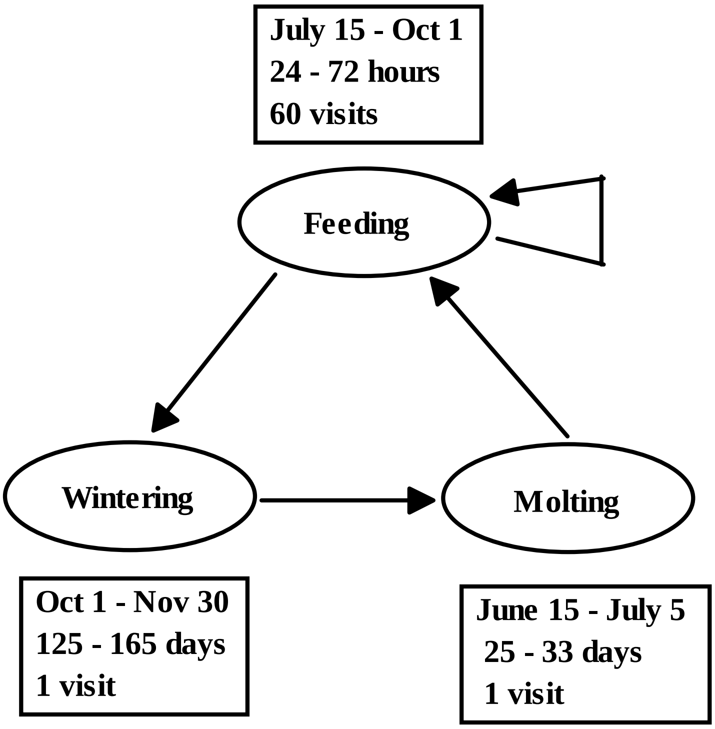

Figure 2: Site graph and temporal site constraints for a reproductive female ringed seal in the Barents sea (Ugland, 1994). The numbers in each box denote the visit season, length and frequency, respectively.

As shown in Figure 2, members of the *reproductive female* class can arrive at wintering/hibernating sites anytime between October 1 and November 30. They can only visit wintering sites once per season, but their visit duration ranges from 125 to 165 days. In contrast, the feeding sites are visited approximately once a day for the entire feeding season, with each visit lasting between 24 and 72 hours.

Note that the structure of the site graph as well as the temporal constraints upon the nodes are user-defined aspects of a MIGMOD scenario, entered via pop-up forms in MIGMOD's graphical interface. This permits a wide variety of migratory patterns that may vary significantly across animal classes, thereby capturing the heterogeneous movements of an entire population.

To keep track of its own migrating history, each animal object maintains a *travel history,* which records the number of previous visits to each site type. As detailed below, MIGMOD integrates all three aspects of the travel itinerary when computing animal movements.

III. **Migration Algorithm**

With the exception of dependent animal particles (e.g. newborns), all MIGMOD animals spend the entire simulation alternating between two activities: 1) randomly wandering about a particular site at their class-specific wandering/searching speed, or 2) migrating from one site to the next at their class-specific migrating speed. The migration algorithm uses the habitat-grid, site-habitat map, tenacity grids, and travel itinerary to stochastically determine both 1) *when* an animal object should stop wandering and start migrating, and 2) *where* it should migrate.

III.1 Object-Oriented Aspects

Below, the pseudocode for *animal::makemove* (i.e., the *makemove* method for the general *animal* class) illustrates the general activity of most animal particles at each time step.

```
animal::makemove()
  save_location()
  if [migrating?() & arrive_goal_site?()] then register_at_goal_site()
     else if [wandering?()] then consider_departing()
  if [wandering?()] wander_one_timestep()
        else migrate_one_timestep()
```

The *makemove* method is inherited by most subclasses in the hierarchy. Only subclasses that have *dependent* or *reproductive* as an ancestor class will have makemove code that differs from *animal::makemove*. Hence, generic *birds,* *immature male seals, and mature female bears* all use the same *makemove* code. They also inherit the same *migrating?, wandering?, arrive_goal_site?* and *consider_departing* methods. However, they typically differ in their *migrate_one_timestep* routines, since, as discussed below, the migratory movement algorithms differ across animal types.

The elegance of OOP becomes apparent when we examine the semantics for "applying" a method to an object, which is analogous to calling a subroutine in other languages. To move any *animal* particle P of any class C, MIGMOD applies the *makemove* method to P, but **without any reference to C**!! The compilers for OOP languages insert extra code for determining C from P. Once C is known, the appropriate *makemove* method is called. Since none of this "dispatching" activity appears in the source code, object-oriented programmers can abstract behaviors into truly general methods that make no reference to specific details. Furthermore, if a new subclass S is later added to the system, one can simply attach it under the appropriate parent class(es) of the hierarchy and define any additional properties or methods. If a local *makemove* method is defined for S, no changes to the methods of other classes are required.

In contrast, a non-OOP approach to the multi-class migration problem would require that all dispatching code be written at the source level. Hence, "general" routines would be huge, containing large "if ...then....else" or "case" constructs that keyed on the (over 50) specific subclass types.

III.2 Stochastic Movement Routines

Intrasite wandering movements are accomplished by simple random-movement routines, which can either restrict the animal to remain within the site, or allow it to wander into neighboring habitat cells of the same site type. The latter constraint is most common, while the former is useful for private sites such as the den of a polar bear. For intersite movements, the algorithm for flying animals is trivial: MIGMOD simply increments the location with a velocity vector at each time step. However, movements on land/ice or in water require a bit more calculation. There, a weighting function is applied to the 8 neighboring grid cells of the animal's current cell. The weights are based on factors such as the directional deviation of a cell from the destination site, the cell's ice coverage, or the cell's water depth. A stochastic choice among the weighted cells yields the animals direction of movement. This enables bears to follow the ice edge and seals to swim around intervening land masses.

Although the above movement calculations are nontrivial for swimming and walking animals, the greatest computational resources are devoted to **determining** the destination sites for intersite movements. When an animal's visitation time at a site has expired, the choice of "where to go next" is eminent. There are two nonindependent aspects to this choice: (a) pick a site type (e.g. nesting, feeding, molting, etc.), and (b) pick a particular habitat cell that corresponds to that type.

The former choice is based on the possible intersite connections in the site graph. For example, from a feeding site, it may be possible to go back to the nest or to begin emigrating from the environment, but not to go to a molting site. Each site type's temporal constraints along with the travel history of the animal come into play. Essentially, for each site type that is legally connected to the animal's current site type, the following weighting function is applied:

$$W = ((f_v - n_v) / f_v) * (t_a - t_1) / (t_2 - t_1) \quad \text{where: } t_1 \le t_a \le t_2 \text{ and } f_v \ge n_v \qquad (1)$$

$$W = 0 \qquad \text{Otherwise}$$

f_v = visit frequency for the site type, expressed in average number of visits per "season" (i.e. the interval between t_2 and t_1).
n_v = number of visits that this animal has already made to sites of this type.
t_1 = earliest visit time (time of arrival) for this site type.
t_2 = latest visit time (time of arrival) for this site type.
t_a = estimated arrival time for this animal at a randomly-chosen site, S, of this type, based on the animal's current location and class-specific migration speed.

This weighting function gives preference to those site types for which the current number of visits by the given animal is well below the expected number (i.e, $(f_v - n_v) / f_v$ is close to 1). The ratio of time intervals in equation 1 accounts for the time remaining to make visits to sites of that type. As this ratio approaches 1, it becomes more urgent to visit sites of this type. If all site types receive a weight of 0, then the $f_v \ge n_v$ condition is relaxed and all sites are evaluated again.

As implied by the definition of t_a, the choice of site type T and site S cannot occur independently: MIGMOD needs a representative site from which to base the calculation of t_a. MIGMOD makes several stochastic choices (from the set of all sites of type T) in search of an S whose corresponding t_a satisfies $t_1 \le t_a \le t_2$. This stochastic selection is weighted in accordance with the tenacity information, wherein the densities in the cells of the tenacity grids are normalized to form visitation probabilities for each cell. In the absence of a tenacity grid for a given site type, other factors such as ice coverage can be used to bias site selection, so that, for example, birds can feed primarily along the ice edge.

When an animal arrives at a site, it receives a visit duration that is randomly chosen from the (uniformly distributed) visitation interval. When that duration expires, the animal must find another site to visit. The site-selection algorithm and random choice of visitation duration lead to distributions of visited sites, arrival times and durations that closely mirror the temporal site constraints.

IV. Biological impact assessment with MIGMOD

To estimate the effects of an oil/chemical release upon a migrating population, MIGMOD inputs a file describing the time-tagged locations and sizes of oil/chemical spillets. For each timestep, the system compares spillet and animal trajectories and tallies "oilhit" incidents when the two intersect in space and time. The duration of the intersection, also recorded, and the vertical thickness of the spillet are used to compute an exposure area for the encounter. The total exposure area (see column 5 of Table 4.1) gives an indication of the extent of the exposure for each age/sex class.

A MIGMOD biological assessment produces three main types of information:

1. The locations and exposures of the simulated animals at periodic intervals (i.e., every few time steps).
2. The locations of all exposure incidents.
3. The final exposure statistics for each age/sex class.

Information of types 1 and 2 is used to graphically display animal movements and exposure incidents, while the statistical exposure summaries can be tabularly displayed and/or used as input to a population dynamics model.

V. Sample MIGMOD scenarios: Ringed seals of the Barents sea

This section illustrates both the generation of migratory patterns and biological impact assessment for two ringed seal scenarios. In the Barents sea, ringed seals appear predominantly in northern areas around the Svalbard archipelago (Ugland, 1994; Mehlum, 1992). In winter and spring, adult seals inhabit the inner reaches of ice-covered fjords, while immature seals prefer the icy coastal areas but are also found within the fjords. Pups are born between mid-March and mid-April. When the approximately 2-month weaning period ends, the adults swim out of the fjords in search of ice-covered molting areas. Males begin to molt in mid-June, while females may not arrive until early July. Meanwhile, the newly independent pups and the older immature seals head toward the ice edge for feeding. After molting, adults also frequent the ice edge. All animals feed along the ice edge until late autumn, when the fjords begin to freeze over. The adults then head back to the inner fjords, while the immature seals migrate toward the coastal regions.

To perform a 1-year MIGMOD simulation of a ringed seal population along eastern Svalbard and Bear Island, the habitat grid of Figure 2 along with a set of site-habitat associations are used to define the areas where seals can appear. Four age/sex classes are simulated: reproductive males, reproductive females, generic dependents, and generic immature seals. The site-habitat associations and tenacity grids
help to manifest the age/sex-specific spatial distributions discussed above. For example, wintering sites for generic immature seals map to both seaward and landward subtidal rock bottom habitats. A tenacity grid weights the seaward (coastal) habitats higher than the landward (fjord) habitats. The wintering sites for mature seals only map to the landward subtidal rock bottom habitats, and the tenacity grid is evenly weighted over those habitat cells.

Temporal site constraints similar to those described above are entered for the various age/sex- specific site types. This insures, for example, that males begin molting before females. Information concerning birth and development restricts the pupping period to the mid-March to mid-April interval, prevents the pups from moving any distance from the birth site for the first two weeks, and enables the pups to become independent after 60 days, at which time they join the ranks of the generic immature seals and head for the ice-edge feeding areas.

A simulated oil spill of 1000 tons medium crude is released on April 1 from a point near the ice edge northwest of Bear Island. The southwestly winds cause the oil to drift north toward Svalbard. The oil evaporates and disperses very slowly due to the cold temperature (-10 C in air) and extensive ice cover. Hence, more than 2 weeks after its release, the spill still poses a threat to animals near Storfjorden (the large southward-opening fjord that nearly divides Svalbard).

For the seal scenarios, MIGMOD is initialized with 300 animals, 100 for each of the three independent age/sex classes, while dependents are born during the simulation. The simulation begins in January, with mature males and females wintering in the fjords, while immature generic seals appear along the coast. Pupping begins in mid-March while all seals are still at their wintering locations. In MIGMOD, each female reproductive animal (dark circle) produces a single generic dependent animal (white dot), which remains in the fjord with the mother during the 60-day weaning period, as shown in Figure 3. This figure also indicates that several animals encountered oil spillets.

By mid-June, most generic dependents have become generic immature seals (light half circles). These immature animals head toward the ice edge in early June. Mature males also head for the ice in order to molt. The females begin molting a few weeks later. In Figure 4, which uses the ice grid for background, note that most seals reside in the 50-60% ice-cover area. MIGMOD uses 50% as the "ice edge" concentration and assigns weights to feeding sites based on their deviation from this "optimal" ice value.

Age/Sex Animal Class	**# Animals Simulated**	**% Animals Exposed**	**Total Exposure Time (hours)**	**Total Exposure Area (m**2)**
Reproductive Female	100	2 (3)	2.6 (0.056)	.35 (0.1)
Reproductive Male	100	3 (7)	7.76 (0.161)	.31 (0.44)
Generic Immature	200	1.5 (8.5)	3.48 (0.33)	.43 (0.88)

Table 1 Exposure data for the ringed seal simulation with an April 1 and June 10 (in parentheses) oil spill near Bear Island.

As the winter sets in mid October, the simulated seals begin swimming back toward the fjords and coastal waters. Since MIGMOD simulates birth and development but not predation, the number of immature seals (light half circles) effectively doubles during the course of the simulation. However, MIGMOD does not model the transition from immature to mature animals; population dynamics models, which can be coupled to MIGMOD, should handle that task.

As this data and Figures 3-4 indicate, the impact is rather local. Since the simulated seals were "wandering about at wintering sites" when the oil spillets reached Svalbard, only those individuals with wintering sites along the oil trajectory were affected. The impact would have been greater if the release had occurred in early June, when a large percentage of the seals were moving out of Storfjorden (the large southward-opening fjord that nearly divides Svalbard) and toward the ice edge. Table 1 also indicates the impacts due to a June 10 release of 1000 tons medium crude from the same location, using the same wind file.

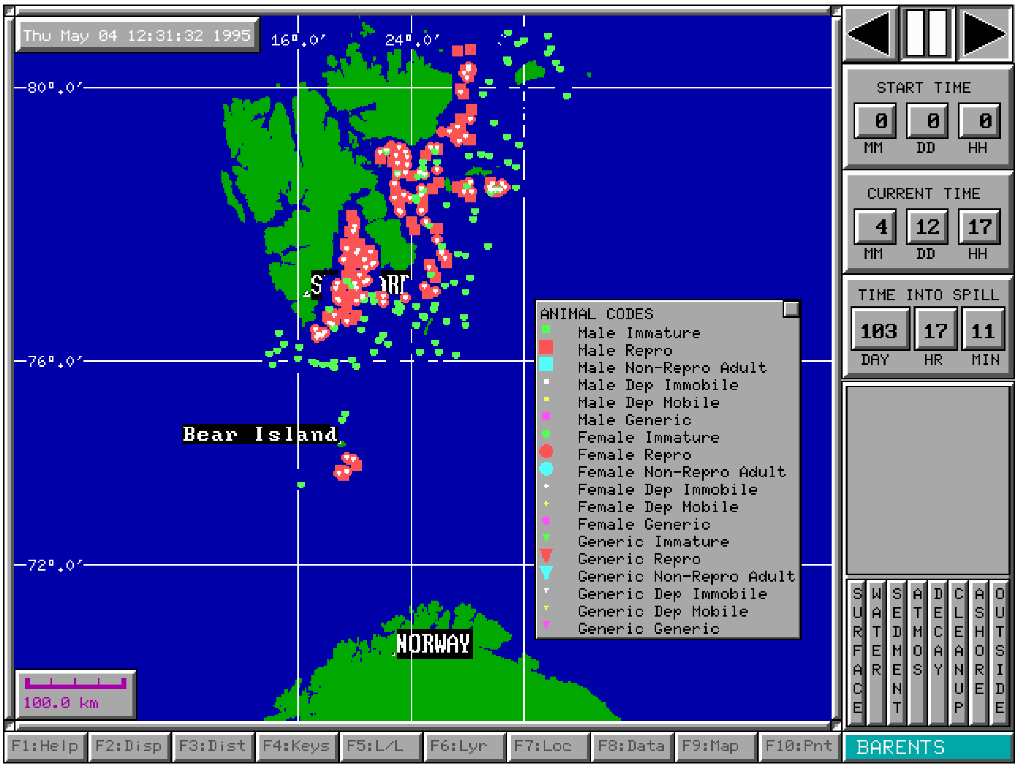

Figure 3. *By mid April, most reproductive females (dark circles) have given birth to a single pup (small white dot), which they wean in the fjords. Mature males (dark squares) winter in the fjords, while generic immature seals (light half circles) stay in coastal waters.*

Comparing the 2 release scenarios in Table 4.1, note that the number of exposed animals increases from the April to the June spill. The increase is small for reproductive females, who do not begin the migration to molting sites until early July. However, this migration is well underway for reproductive males and immature generics when the oil spillets reach Storfjorden. Hence, more of these animal types encounter oil in June than in April. In the April scenario, animals met oil while wandering about at a site. Since MIGMOD does not model the animals avoidance of oil, the animals were exposed for relatively long periods of time. Conversely, the reproductive male and immature generic animals in the June scenario were moving at a migratory speed when they met oil. Since the modelled migratory speed of 5 km/hr exceeds the drift velocity of the oil patches (e.g. 30 cm/sec in a 10 m/sec wind), the exposure times in the June scenario were shorter than in the April scenario, where spill drift velocities largely governed the exposure times.

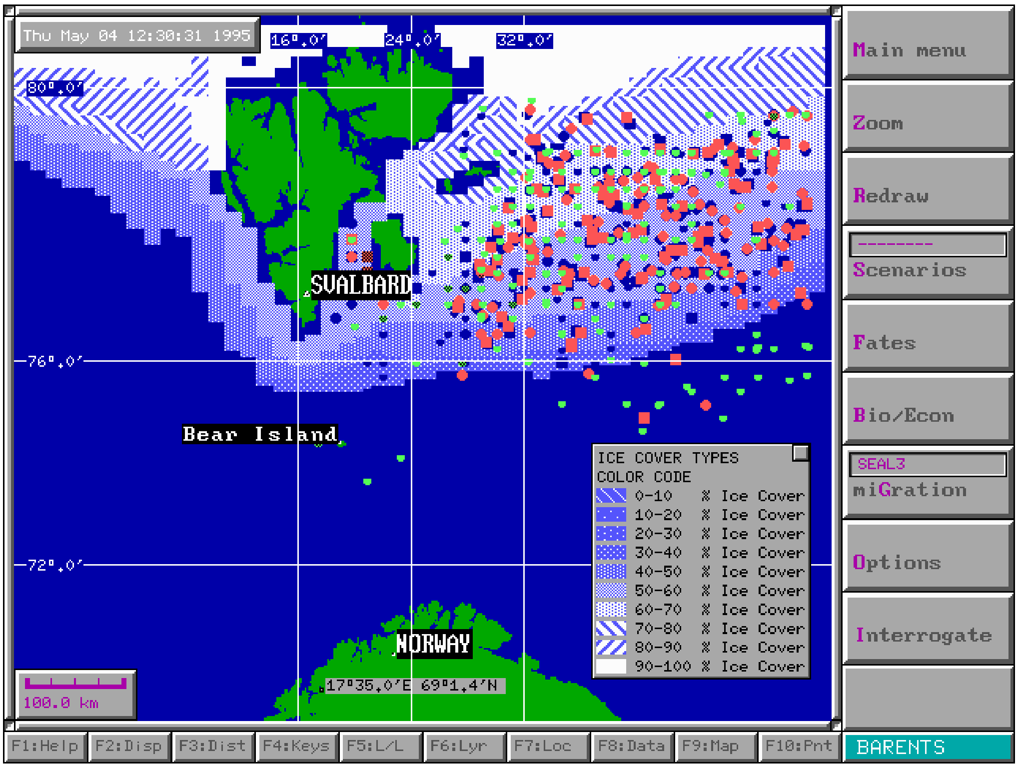

Figure 4. *In mid June, the generic immature and male reproductive seals migrate toward the ice edge east of Svalbard. Reproductive females lag a few weeks behind. By early July, all classes have arrived at the ice edge. Animals hit by oil are indicated by checkered patterns.*

VI. Discussion and Future Research

MIGMOD exploits the data encapsulation, class hierarchies and inheritance of the object-oriented paradigm to efficiently handle many animal species, subclasses and idiosyncratic migrational behaviors. This approach enables MIGMOD to manage the complexities of heterogeneous population modelling, thus providing a general-purpose, multi-species, multi-class migration model.

The task of migratory simulation illustrates an interesting trade-off between causal or first principle modelling approaches and more empirical methods. Although a wide variety of migratory methods and behaviors exist across the biosphere, there is no need (for the purposes of acute biological impact assessment) to model detailed physical mechanisms such as hitching, guiding or navigation (Dusenbery, 1992; Balchen, 1992), when the statistical large-scale behaviors are already known: biologists often have good data on when species arrive and depart from certain key sites. Given this temporal and spatial information, our problem becomes one of stochastic scheduling with a few appeals to the ambient physical data, for example, to keep simulated polar bears walking along the ice edge.

The computational complexity spared by the reduction in sensory modelling is quickly replaced by the demands of MIGMOD's stochastic processes. The weighted random selection of one site among thousands, performed every few (simulation) hours for every animal, greatly increases computational time. Future MIGMOD research will focus on streamlining these stochastic methods and replacing some of them with efficient models of relevant causal migratory mechanisms such as hunger-driven initiations of feeding forays.

Another future extension involves the coupling of MIGMOD to a differential-equation based population dynamics model. Briefly, field data would provide an initial population distribution to MIGMOD, which would then create the appropriate ratio of age/sex class objects. Births, maturations and pollutant exposure events during the run would lead to an updated population distribution, with exposures above a given threshold entailing death. The population dynamics model would then run for several years on this new distribution to indicate the long-term populational effects of an acute (i.e., short term) environmental disturbance.

MIGMOD is currently restricted to running one animal species at a time. Although independent MIGMOD scenarios for birds, ringed seals, and polar bears have been tested (Downing, Reed & French, 1994), a model of multiple interacting species would enable a more comprehensive migratory ecosystem analysis. For example, seals and birds would follow schools of fish, while polar bears would follow the seals. An integrated ecosystem model of this type could greatly improve impact assessment, particularly in cases where the usual site and visitation date required by MIGMOD is unavailable for several of the relevant species. This integrated model would incorporate both Lagrangian and Eulerian modelling methods, with the Lagrangian movements of prey individuals used to compute prey concentrations (in Eulerian grid cells), which would then govern the Lagrangian movements of their predators, which in turn would alter prey concentrations via feeding.

This interacting species approach would push MIGMOD further in the direction of causal-based ecosystem models and away from its empirical data-driven roots. Since the field-data demands of MIGMOD are already so great, future versions will require either more causal mechanisms or a direct connection to an extensive wildlife database. Elements of both solutions are currently under investigation.

In general, the ability of OOP to span the generality/specificity spectrum should help address key issues of pattern and scale in ecological modelling (Levin, 1992). OOP models such as MIGMOD can cover a reasonable portion of Levin's desired "suite of models of varying levels of complexity" under one roof. This added breadth at the local/individual level facilitates a more robust investigation of the contributions of local behaviors to global patterns.

References

BALCHEN, J.G., 1994: Modelling of Fish Behaviour. Proc. Environmental Modelling Seminar, August 17-19, 1994, Trondheim, Norway. p.287-303.

DOWNING, K., M. REED, and D. FRENCH, 1994: Biological Effects of Potential Oil Spills at the Ice Edge in the Barents Sea. IKU Report # 43.2171.00/01/94.

DUSENBERY, D.B., 1992: Sensory Ecology: How Organisms Acquire and Respond to Information. W.H. Freeman and Company, New York.

FORD, R.G., J. WEINS, D. HEINEMANN, G. HUNT, 1982: Modelling the sensitivity of colonially breeding marine birds to oil spills: guillemot and kittiwake populations of the Pribilof Islands, Bering sea. J. Applied Ecology 19: 1-31.

HOGEWEG, P. and B. HESPER, 1990. Individual-oriented modelling in ecology. Mathematical Computational Modelling, Vol. 13, No. 6, pp. 83 - 90.

LEVIN, S., 1992. The problem of pattern and scale in ecology. Ecology, 73(6), pp. 1943-1967.

MEHLUM, F., 1992: Birds and mammals of Svalbard. The Norwegian Polar Research Institute.

REED, M., 1983: A multidimensional continuum model of fish behavior. Ecological Modelling. 20, pp. 311-322.

REED, M., J.G. BALCHEN, 1982: A multidimensional continuum model of fish population dynamics and behaviour: application to the Barents Sea capelin (Mellotus vilosus). Modelling Identification and Control 3(2):65-109.

UGLAND, K., 1994: Personal communication.