

Modeling Agents and their Environment

James J. Odell
James Odell Associates
3646 W. Huron River Dr.
Ann Arbor, MI 48103
Tel: +1 (734) 994-0833
jodell@compuserve.com
www.jamesodell.com

H. Van Dyke Parunak
ERIM
P.O. Box 134001
Ann Arbor, MI 48113-4001
Tel: +1 (734) 623-2509
vparunak@erim.org
www.erim.org/~vparunak

Mitch Fleischer
ERIM
P.O. Box 134001
Ann Arbor, MI 48113-4001
Tel: +1 (734) 623-2517
mfleischer@erim.org
www.erim.org/~mfleischer

Sven Brueckner
ERIM
P.O. Box 134001
Ann Arbor, MI 48113-4001
Tel: +1 (734) 623-2529
sbrueckner@erim.org
www.erim.org/~sbrueckner

Abstract. Without an environment, an agent is effectively useless. Cut off from the rest of its world, the agent can neither sense nor act. An *environment* provides the conditions under which an entity (agent or object) can exist. It defines the properties of the world in which an agent will function. Designing effective agents requires careful consideration of both the physical and communicational aspects of their environment. Two issues exist for understanding environments:

1. Every agent has an environment, no matter what the agent's philosophy or architecture is.
2. Being aware of the agent's environment enables its designer to get more powerful interaction via architecture-dependent means.

1. Introduction

Agents need to operate and exist within an environment. Figure 1 illustrates a common view that agents perceive their environment through sensors as well as effect actions on it. [Pfeifer, 1999; Weiss, 1999; Russell, 1995] For example, a Stock agent can receive an event indicating that quantities of a particular part are low. The agent then decides whether more parts need to be ordered and, if so, put out a general call-for-proposal so that interested vendors can reply. When proposals arrive, the Stock agent will choose and notify the winning vendor. This model implies that agents interact *via* an environment. Even direct communications (such as vendor notification) must occur through some medium. In other words, the environment provides the appropriate conditions that enable interaction among agents. This insight, largely overlooked in the design of purely electronic agents, is particularly critical for managing agents that are situated in the physical world.

2. What is an Environment?

One of the key properties of agents is their autonomy. However, autonomy is not an all-or-nothing issue. Practically speaking, agents can neither be totally free of external dependencies nor completely reliant on them. They always depend on external factors to some degree.

An *environment* provides the conditions under which an entity (agent or object) exists.

In other words, it defines the properties of the world in which an agent can and does function. An agent's environment, then, consists not only of all the other entities in its environment, but also those principles and processes under which the agents exist and communicate. Designing effective agents requires careful consideration of all of these factors when designing their environment.

A canonical example of agents situated in an environment is an ant colony. Ants interact with one another largely through chemicals, called “pheromones,” that they deposit in the environment and then sense to guide their actions. Numerous individual interactions yield the emergent development of paths through the environment. However, the environment is more than just a communication channel. Agents depend both on tangible, physical support and on other agents. Two aspects, then, are critical for agent environments (and the formation of paths): the physical and the communicational.




Figure 1: Agents interact with and through their environment.

3. Physical Environment

The particular kind of environment that biological agents (animals and plants) require for survival is referred to as their ecological niche. Edward O. Wilson defines *ecological niche* as: “The range of each environmental variable such as temperature, humidity, and food items, within which a species can exist and reproduce.” [Wilson, 1975] While artificial agents can have different requirements for survival, they still require an ecological niche, or physical environment, to support them.

The *physical environment* provides those principles and processes that govern and support a population of entities.

Principles

The laws of physics provide us with the fundamental truths that are essential to the world in which we live. For example, a physicist could use the study of particle dynamics to describe the causes for motion and the way in which bodies influence each other. For such descriptions, we obtain principles such as the conservation of energy, gravity, sound waves, and fluid dynamics. In Karl Sims’ agents, the same principles apply because his “creatures” were bred to swim, run, and fly in a world whose laws of physics are almost identical to ours. [Sims, 1994a, 1994b] In contrast, the ant’s environment has its own particle dynamics. For example, ants may only move from one place to an adjacent place; no two ants may occupy the same place at the same time; and yet pheromones may be aggregated when separate ants deposit them at the same place. The concepts of diffusion and evaporation are also part of the agent environment. This makes it

possible for pheromones to spread to neighboring places as well as evaporate over time. Similarly, a statement of fundamental qualities is also required for agent environments. Here, each agent-based system must identify and define those fundamental truths forming the ground of its system.

For agents, principles of the physical environment can be thought of as laws, rules, constraints, and policies that govern and support the physical existence of agents and objects. Basic characteristics for an agent environment can include [Weiss, 1999; Russell, 1995]:

- **Accessibility**- To what extent is the environment known and available to the agent? An environment is effectively accessible if the agent can access the environmental state relevant to the agent's choice of action. Another consideration is whether the available resources are ample or restricted.
- **Determinism**- To what extent can the agent predict events in the environment? The environment is deterministic when the next state of the environment can be determined by the current state and the actions selected by the agents.
- **Diversity**- How homogeneous or heterogeneous are the entities in the environment?
- **Controllability**- To what extent can the agent modify its environment?
- **Volatility**- How much can the environment change while the agent is deliberating?
- **Temporality**- Is time divided in a clearly defined manner? For example, do actions occur continuously or discrete time steps or episodes?
- **Locality**- Does the agent have a distinct location in the environment which may or may not be the same as the location of other agents sharing the same environment. Or, are all agents virtually collocated? Also, how is a particular locality expressed (e.g., coordinate system, distance metrics, relative positioning)?

Processes

Formally, an environment can be expressed as a two-tuple [Parunak, 1996]:

$$Environment = \langle State_e, Process_e \rangle$$

Where, $State_e$ is a set of values that completely define the environment. The structure, domains, and variability of these values are not constrained by this definition, and differences in these features are responsible for much of the interesting variation among different kinds of environments. The state also includes the agents and objects within the environment. $Process_e$ is an autonomously executing mapping that changes the environment's state, $State_e$. "Autonomously executing" means that the process runs without being invoked from any outside entity. In computational terms, an environment has its own virtual CPU. The important feature of this definition of environment is that the environment itself is active. It has its own process that can change its state—which includes the agents and objects within the environment—independently of the actions of its embedded agents.¹

In an agent environment, the primary purpose of these processes is to implement the environmental principles. For example, the gravitational field is a principle that can be implemented with a process that attracts entities in a prescribed manner. In other words, the falling of an apple to earth can be regarded as the process of gravity in action.


¹ The exact nature of the coupling between agents and their environment depends on how state and process are modeled in each: as a discrete-event or time-based dynamical system. The former involves a discrete state with a symbol-manipulation processing style; the later, a continuous state with difference or partial differential equations. See [Parunak, 1996] for more details.

In the case of ants, the environment is not a passive conduit for information. Instead, it actively processes pheromones in three ways. It aggregates pheromone deposits from separate ants at the same place (thus realizing a primitive form of information fusion). It evaporates pheromones over time (thus providing a novel form of truth maintenance). Finally, it propagates pheromones to neighboring places (thus disseminating information). Experiments show that these mechanisms are critical to the formation of paths. More generally, environmental activity means that the environment may change even when the agents living in the environment do not take action.

Different physical environments will be required for different kinds of agents—and vice versa. With artificial agents, much more than physics is happening because much of the environment is information intensive. In ant-based environments, the pheromones *are* information. In many defense-related agent systems, the information-intense environment includes satellite telemetry, body- and vehicle-based communications technology, and geographic positioning grids. In agent-based supply chains, information about orders and resources is a major component of the system.

To support the varied information requirements of such agent-based systems, a common processing platform would be useful. This platform would provide a foundation upon which agent applications could build to leverage their own specific environmental requirements. Figure 2 illustrates the primary components required to support a physical environment for agent-based applications—whether the agents are implemented as software, hardware, or a combination of both:


- **Application Support** contains the applications, as well as all management and support services for the entities supported by the environment, such as directory and ontology services, query, mobility, security, and firewalls.
- **Communication and Transportation** packages, routes, verifies, and transmits data required for the application support layer. It provides a general-purpose service that has no application dependencies and the type of data does not matter.
- **Physical Linkage** specifies the physical and electrical characteristics of the bus. Typically, this involves the hardware that converts the characters of a message into electrical signals for transmitted messages and electrical signals into characters for received messages. This can include standard physical interfaces such as controllers, actuators, sensors—as well as road networks and pallets.



The processes for an agent's physical environment may be implemented in either hardware or software; however, at some point (Physical Linkage) the environment must be realized in some material form. For example, CAN (Controller Area Network) has developed hardware for Physical Linkage layer. They have also developed software for the Application Support layer that supports CAN controllers and interface devices [<http://www.omegas.co.uk/CAN/>].

Some work has already been done to define the standard services required for agent-based physical environments. The FIPA (Federation of Intelligent Physical Agents) *Agent Platform* defines an abstract architecture for agent deployment and is summarized in Fig. 3. [FIPA, 1998] The existence of layered protocol such as FIPA and ISO shows that people already have an intuition about the importance of relating agents to the rest of the world.

- **Agent management system (AMS)** can be implemented as a single agent that supervises access to and use of the agent platform. The AMS maintains a directory of logical agent names and their associated transport addresses for an agent platform. The AMS is responsible for managing the lifecycle of the agents on the platform and actions such as authentication, registration, de-registration, search, and mobility requests.
- **Agent platform security manager (APSM)** is responsible for maintaining security policies for the platform and infrastructure. The APSM is responsible for run-time activities, such as communications, transport-level security, and audit trails. Security cannot be guaranteed unless, at a minimum, all communication between agents is carried out through the APSM.
- **Agent platform communication channel** provides a path for basic interchange between agents, agent services, AMS, and other agent platforms. It must at least support IIOP. Agents can reach agents on any number of other platforms through the Agent Communication Channel. Ways of communicating include using blackboard or message-based communication; point-to-point, multicast, or broadcast; push or pull; and synchronous or asynchronous.



In spite of the acronym, the FIPA architecture focuses almost entirely on the electronic environment, and does not address the physical environment. As such, it does not address the real potential of an active environment to provide emergent system-level behavior. As stated earlier, every agent has an environment. However, such environment can be consciously used in special ways to get more powerful interaction.

A standard that does address the physical environment is the ISO/OSI model, depicted in Fig. 4.² This model describes how communications should occur between computers on any network, and has been adopted as a general "open" network communication standard. In principle, anything that conforms to the standard can communicate, electronically, with anything else that conforms to the standard.

Population

An environment is an inhabited place; i.e., it is populated. An agent's environment might or might not contain other entities, and it might be open or closed. An environment's population is the totality of entities under its consideration. For the environment of the canonical software ant, this population would consist of food, pheromones, and other ants. For a real-world ant, it would also include earth, twigs, trees, and picnics. For a stock agent in a supply network, it would include physical inventory, road and rail networks, packaging conventions, and so on.

4. Communication Environment

In *individualist* agent environments, agents are viewed as independent entities; whereas in *collectivist* environments, agents are viewed as interdependent. While an agent can operate by alone, the increasing interconnections and networking require a different kind of agent—one that can communicate effectively with other agents. A communication environment provides two things. First, it provides the principles and processes that govern and support the exchange of ideas, knowledge, information, and data. Second, it provides those functions and structures that are commonly employed to enhance communication, such as roles, groups, and the interaction protocols between roles and groups. In short:


The *communication environment* provides those principles, processes, and structures that enable an infrastructure for agents to convey information.

Communication

Basically, communication is the conveyance of information from one entity to another. The nature of this transfer can range from the simple to the complex. For example, a satellite could periodically send one bit to inform ground control that it is still functioning correctly; in contrast, the information exchanged within the US Senate to negotiate tax cuts can appear quite chaotic. In contrast, broadcasts such as television commercials do not necessarily result in communication. A signal may go out, but if you are not listening or watching, how can the commercial convey information?


Figure 5 illustrates the difference between transmission and communication. In Fig. 5(a), neither agent has any transmission activity. Figure 5(b) indicates that the agent on the left transmitted information through the environment, but was not received by the other agent. Communication, however, requires that the information transmitted by one agent results in a state change of another (Fig. 5(c)). In the case of television commercials, perceiving its transmission means that your senses have at least detected it. The perception could involve you buying the advertised goods, throwing a shoe at the television screen, or simply choosing to do nothing. Either way, communication has occurred because the act of sensing and deciding involves a state change by the receiver.

² Guy Genilloud, Guy has proposed a flexible translation for linking FIPA to OSI via CORBA in [Genilloud, 1997].



Interaction

Proving that communication has occurred, however, requires us to know that the inner state of the receiving agent has in fact changed. We are not advocating that the communication environment possess such mentalistic knowledge—only that such an environment be present so that transmission and communication can occur. However, knowing that a transmission was received can be important to the sending agent. One useful way to determine if communication has occurred is when an interaction results. Figure 6(a) depicts one agent communicating with another. Here, the other agent responds, but the original agent does not receive the responding transmission. (The original agent, then, cannot know for a certainty whether communication




occurred.) In other words, there was no interaction between the two agents. *Interaction*³ requires two-way communication (i.e., a reciprocal effect), as illustrated in Fig. 6(b). Interaction, then, not only defines exchange of information, it confirms that the original transmission was in fact received by the other agent. In other words, the original agent can infer that its transmission was communicated to the other agent as soon as a response is received—even if the response communicates only that the responder did not understand the original message.

Social environment

In agent-based systems, communication and interaction are commonly employed together. Furthermore, agent-based communication can even involve patterns of interaction, or *interaction protocols*. From simply requesting the price of a product to conducting elaborate contract-bidding activities require that some agree-upon approach be in place to facilitate interactive communication—without which the conveyance of information could easily result useless Babel. Such a situation could be considered *social*.

A *social environment* is a communication environment in which agents interact in a coordinated⁴ manner.

As illustrated in Fig. 7, the social environment is a subset the communication environment. In other words, not all communication is social (as defined above), but all social activity requires communication.



³ The action or influence of agents on each other; i.e., having a reciprocal effect.

⁴ Coordination - the act of agents performing activities in the context of objects and other agents in a shared environment.

Principles

Communication principles provide us with the fundamentals that are essential for interactions, customs, norms, values, commitments, dependencies, and so on, that constitute an agent society. The canonical ant's communication environment is simple: all communications between ants are via pheromones. Here, the communication involves a two-step process: the ant deposits pheromones that act as information for other ants, while the "other ants" query the local environment for the presence of pheromones. In short, these ants participate in a social communication by way of environmental substances, rather than directly with each other. Additionally, some ant societies have multiple kinds of pheromones: one for exploring and one when returning with returning food. The net effect is informing a society of ants about how to find food or home—whichever is useful for any given ant. If an ant is foraging, information about where to find food is useful; if the ant has food, directions on how to get back the colony would be useful. Such interaction is social because it provides the ants with an infrastructure for the colony because it produces coordination among the ants.

Supply-chain agents can have elaborate collaborative protocols acquiring and delivering goods and services along value-adding chains. Defense-related protocols, require different interaction policies at different command levels. Both direct and indirect interaction can be employed as interaction strategies in these applications.

In rich multiagent societies (MAS), several principles are required to facilitate the communication environment.

- **Communication language-** Agents communicate to understand and be understood. The formal study of communication has three aspects: syntax, semantics, and pragmatics. Agent-based social environments must define the principles required to address these aspects. Additionally, it must define the types of messages that will be employed (e.g., assertions, queries, replies, requests, and denials) and the ontology. Some of the common agent communication languages (ACL) languages include FIPA ACL, and KQML.
- **Interaction protocols-** An agent interaction protocol (AIP) describes a communication pattern as an allowed sequence of messages between entities and the constraints on the content of those messages [Odell, 2000a and 2000b]. Examples of AIPs include the contract net protocol, Dutch auction protocol, and publish/subscribe protocol. FIPA has standardized more than a dozen AIPs [FIPA, 2000].
- **Coordination strategies -** Agents communicate to achieve their goals and the goals of the social group in which they participate. Cooperation, competition, planning, and negotiation are common principles used to perform activities in a shared environment. AIPs can be associated with each of these strategies.
- **Social policies-** The permissions and obligations that dictate acceptable social behavior. They include being able to apply and enforce these policies across distributed agents and systems. The general focus here is on the application and management of policies on agents and groups of agents—not the detailed management of agent lifecycles and areas currently addressed by FIPA agent management specifications. Other considerations for social policy can involve:
 - Implicit vs explicit rules; not all rules are specified in advance: i.e., learning what the rules are or adjusting to a change in rules, emergence of rules, unconscious rules, when tradition becomes a social norm, or policy.
 - Different levels: of influence/power (e.g., the ROI on obeying or violating a rule, strength and "evaporation" of rule; rules don't always stick around, rules as memes, language use.)
- **Culture-** a set of values, beliefs, desires, intentions, trust, morality. These can determine the characteristics of the above. FIPA vs. KQML cultural differences; English vs. other

different-culture language (e.g., Navajo). Culture also affects language, interaction protocol, and social policies (implicit & explicit).

Processes

An agent's communication environment provides processes that enable agents to interact productively.⁵ In particular, it must provide:

- **Interaction management**- managing the interactions among entities to ensure that they are adhering to the selected agent interaction protocol (AIP). AIP adherence can be maintained by the agents participating in the protocol, so that the environment does not need to be involved. However, trusting that each agent can and will adhere to and ensure correct AIP interaction may not be enough to ensure social order. An environment-level control can be implemented as an AIP-manager agent. Did you get what you wanted/needed/expected;
- **Language processing and policing** – where the language parses correctly, it parses correctly but is wrong (evidence or contradictory), or is correct but inappropriate within the agent's context.
- **Coordination strategy services**
 - **Directory service**- locating agents can be supported by white-page (individual), yellow-page (industry), or green-page (offered services) methods. In the physical environment, this directory is used to provide information about where the agent is physically; in the social environment, it provides information about an agent's role or the services that it can provide.
 - **Mediation services**- acting through an intermediate agency. Specialized agents could be established in the environment to act as a communication's intermediary for activities such as transaction management or ontology translation. Environment-level mediation can be implemented using specialized agents.
- **Policy enforcement service**- control of the agent by its environment or social group. The range of possible mechanisms for enforcing policy mechanisms can range from social sanctions to a complete withdrawal of supporting services for the non-conforming agent.
- **Social differentiation**- the process whereby a group or community becomes separate or distinct. To ensure success, groups will institutionalize and employ roles for their members. An agent can play multiple roles in multiple groups.
- **Social order**- the production of a structure of relationships among social agents [Castelfranchi, 2000]. Social order can be the result of formal policies as well as emerge via self-organizing mechanisms. The later is a emergent social pattern of its own, such as the stock market. The former has to do with managing the conditions of an agent society as a whole employing a non-accidental and non-chaotic pattern of interactions. For example, auctions employ strict social patterns. Such a mechanism can be employed to control undesirable emergent patterns that need to be remedied. For example, when stock prices rise or fall by too many points in a session, trading curbs are triggered.

Content

As mentioned earlier, a physical environment consists of all those entities in the physical environment. In contrast, an agent's social environment consists of

- those social units in which that agent participates,

⁵ The agent communication channels are defined as part of the physical environment. The communication environment uses those channels to convey information.

- the roles that are employed for social interaction,
- all the other members who play roles in these social units.

Each social unit, or *group*, is a set of agents associated together by some common interest or purpose. There are three reasons for creating groups.

- *Intragroup associations*- Groups are commonly formed to foster or support the interaction of those agents *within* the group. Here, the group provides a place for a limited number of agents to interact among themselves. For example, such agents might wish to exchange information or seek safety in numbers.
- *Group synergy*– Social units can be formed to take advantage of the synergies of its members, resulting in an entity that enables products and processes that are not possible from any single individual. Corporations, unions, and governments are examples of such social units.
- *Intergroup associations*- Social units also serve as an entity with interactive capability. Here, a group is a set of agents that interact with other sets of agents. Recurrent patterns of interaction define roles, and frequently associated roles are usually considered as defining (sub)groups.

A group can be empty if no agents participate in the group; its collection can also contain a single participating agent or multiple agents. Groups have a separate identity within a larger whole and can be composed of agents, as well as other groups⁶. Furthermore, groups can become social actors influencing group processes and outcomes, as well. For example, most business organizations interact with sector groups such as industry, technology, agriculture, and government; and each of these can influence the other as well as consist of their own subgroups. In this way, an agent social environment can be thought of as a society where agents interact in a more or less ordered community.

A *role* is an abstract representation of an agent's function, service, or identification within a group. In other words, each role is a class of agents that participates in pattern of dependencies and interactions in a prescribed manner. A pattern of dependencies is an important component of a role. For example, if agent A is a customer, there must be some agent B on whom A depends for goods and services, while B depends on A for money. For AIPs, roles define which actions are permitted for a certain class of agents. For example, an agent playing the customer role may request goods, but not supply them; the supplier has the opposite requirements. [Parunak, 2001a]

5. Spatial and Temporal Considerations

An agent's environment—physical or social—must occupy both space and time. Agent populations abide and interact, their processes occur, and their environmental principles are defined over that same temporal space. Agent space and time involves the notion of agent *place*, along with two of its primary attributes: *extent*, and *locality*.

Place

Each agent environment can be thought of as a whole or it can be subdivided into discrete *regions*. Regions partition the agent's physical environment into smaller physical units—where each region may have different or unique characteristics. For example, a grid structure can be defined for the ants so that discrete locations are provided for both the ants and their

⁶ Some debate exists about whether a single agent can be its own group, because each agent can be thought of as having both a social and physical existence. There is another debate about whether or not a group has the status of an agent (holonics vs. AALAADIN).

pheromones, as well as the ability to form pheromone paths. Region definition can also include geographic-based attributes, such as lakes, hills, roads, and structures. In social environments, regions spatially partition the environment into groups and roles. In contrast, temporal space can define unique characteristics for each place in time.

Region specification can include various constraints. For example, in a physical environment we may wish to specify that no two ants may occupy the same place at the same time; yet, we may permit accumulation of multiple units of pheromones. In a social environment, business organizations might be limited to having one person occupying the role of president at any point in time.

Set theoretic distinctions can be made between membership and set. For example, if Agent A belongs to Organization B and Organization B belongs to a federation of organizations C, A does not necessarily belong to C. However, if an Agent A is an element of Set B and Set B is contained by Set C, then Agent A is also an element of Set C.

The region size is determined based on the design granularity: meter-sized places are unrealistic for small ants; micron-sized regions would push the limits of current technology. For example, Pacific Gas & Electric specifies a longitude and latitude within two meters accuracy called a *geocode*. The geocode place size for PG&E, then, is four square meters. In combat examples, a similar grid structure and size is also employed.

In another example, SRI proposes a new top-level Internet domain called .geo. [SRI, 2000] In a .geo system, the Earth would be partitioned into cells based on latitude and longitude. Dedicated servers would hold the data registered to Web sites within its geographical domain, as well as maps and other information. As illustrated in Fig. 8, places can be arranged hierarchically so that search engines could direct queries to one type of server, depending on what the Web user was looking for. The Internet user could then query for cardiac bench surgery in North America or men's clothing stores in Ann Arbor, Michigan. In this way, web user would never need to use unwieldy .com addresses; the geo-enabled search engine translate a geographic location into web sites registered at that location.


Extent

Agent environments must exist in some designated area (or volume) in space and time. Region designations can be expressed in various ways: length/width/height, location points indicating the boundary, memory or disk locations, to and from dates, and so on. The shape of physical space can also be considered here. Social space can expressed in terms the degree of interaction. For example, this could include the number of people you work with, the "degree of separation" between one website and any other website. For example, the environment could be a flat plane or a torus space. In other words, agent environments require an extent that defines its size, shape, and boundaries. Effects of boundary conditions can also be addressed here.

Locality

The ability to locate an entity is an important factor, particularly in agent communication. *Locality* provides the position or situation of a region or entity. Often the locality of the region become an agent's locality. Locality can be addressed in an absolute and relational manner.

Absolute locators are locators that assign a unique address to each agent or region. Simple two-



dimensional grid system employ column-row designations, geospatial systems employ longitude-latitude-altitude designations, and IT systems employ unique identifiers in the form of keys and unique reference IDs.

In contrast, *relational* locality means that an agent's location can be described as relative to another location. For instance in a connected graph, one agent could be related to other agents, which could in turn be related to other agents. Connected graphs such as the web, electric power networks, or networks of colleagues are examples of where entity's location can be described relative to other entities. In a planar environment an agent's relative neighborhood could be based on physical proximity rather than edges between nodes. For example in a simple two-dimensional grid like a checkerboard, one square can be characterized as diagonal to, or to the side of, and so on. This kind of locality is particularly useful when an entity is constrained to interact with the region of the environment that is near it. For example, ant agents may only move from one region to an adjacent region, and their pheromones might "flow" into neighboring regions where pheromone strength lessens the further it travels.

Locality is useful for several reasons. One primary reason is that communicating with an agent requires that the message can actually be delivered to the agent. The sender of a communication may not be required to know where the receiver is physically located, but at some point the communication service must find the receiver to deliver the message. Another reason is to provide location information. For example, a dispatcher agent might need to know the physical location of its various resources to schedule effectively. Lastly, agent movement or interaction




Figure 9: Location-aware computing. [Buderi, 2001]

may be based on, or limited to, physical proximity. For example, an ant agent may only move to or interact with the region that is immediately adjacent to it. In contrast, a flea or grasshopper-style agent may jump multiple squares in a single bound, but is limited to a maximum of five.

Also, it should be noted that since regions are positioned in an environment, the notion of locality applies to them, as well. Furthermore, an agent's locality can be based on the locality of the region it occupies. For example in the .geo example (Fig. 8), each region had an absolute locator within which other entities are contained. In this way, an agent's locality can be defined in a discrete space, instead of locating the agent one large, continuous environment.

Hybrid approaches using absolute and relative locality are also useful. AT&T's "bat" transmitters are a good example of a hybrid approach. "Bats" are small battery-powered ultrasonic transmitters that can be worn on a belt or placed inside objects. They broadcast a uniquely-identifying 48-bit pulse to receivers embedded every 1.5 meters in ceilings as illustrated in Fig. 9. (For example, about 800 are placed around AT&T's three-story lab in Cambridge, England.) Based on the known position of each receiver, the bearer's precise position can be calculated. In other words, the transmitters and receivers have absolute locators: the transmitter has an 48-bit ID and the receivers are coordinate-based. Then, based on relative proximity, the coordinates of the transmitter can be derived from the receivers' coordinates.

Using this location information, zones can be established around objects and people. If a person's zone overlaps an object's zone, the person becomes the temporary owner of the device, be it a workstation, digital camera, telephone or anything else. There is no logging on and everything the user creates—documents, pictures, memos—is automatically stored in the user's personal files. [Baduri, 2001] In other words this technology, known as *location-aware computing*, detects when you're online and what kind of device you're using. Many companies now have development efforts that involve location-based computing: AT&T's Sentient Computing R&D (described above), IBM's Pervasive Computing Division, HP's CoolTown project, the ubiquitous-computing projects at Intel and Xerox. [Want, 2001] Microsoft is another such company with its new HailStorm services platform. When someone tries to get in touch with you, the HailStorm system will detect your network location and level of accessibility: Are you at your desk? In a meeting? In transit? Depending upon the answer, the system will e-mail, page or call you.

6. Designing for single versus partitioned space

As mentioned above, the environment provides those principles and processes under which its population can exist. In implementation terms, this means each environment must become an entity that ensures all environmental principles and process will by properly carried out. Based on the size, complexity, and response requirements, a design decision must be made as to whether there should be

- one centralized environment (i.e. a single region) which oversees all environmental processing;
- a partitioning of the overall space into discrete regions, where each region oversees only the environmental processing within its boundaries; or
- a continuous environment that supports an agent wherever it is. Here, the agent relies on internal location management, instead of external.

Environments treated as a single space would be following a central-control philosophy. In essence, it would require that the space to be omnipotent and omniscient because it must ensure that the physical and social processing is carried out for all entities. With a few agents, this would not be a problem; with many agents, this could easily overwhelm the environment's processing capabilities. Self-managing regions, then, could be a useful design solution for highly populated environments because it partitions the environments processing requirement. Here, the size of each region would need to be determined based on the size, complexity, and response requirements. Hierarchies of places could also be used to accommodate various level of granularity similar to that illustrated in Fig. 8.

A common solution to ensuring the correct space-related behavior is to implement each environment and/or region as an agent. This agent could manage the both physical and communication environment for its designated place based on the global physical and processing requirements. For example, Fig. 10 illustrates having one environment agent associated with the environment. This *environment agent* is also associated with the population of *domain agents* it supports.

When the environment is partitioned into discrete regions, a similar—but more involved—approach is required. Here, each of the regions can be associated with *place agents* that manage





Figure 10: Environment agent supporting multiple domain agents within an environment


environment. The domain agents form another group, where each domain agent is associated with a particular place via its place agent.

Additionally, when the environment is partitioned into multiple place agents, each may have its own local environmental processing overrides and supplements. For example, the properties of water are different from the air and land. An agent in water would be subject to different processes than one walking on land. So, a place agent for a lake would have different properties than an adjacent land agent. In the computational setting models we have constructed of ant systems, we model the environment as a network of agents, which we term *place agents* (or *places*, for short). Because they are agents, places can run processes, thus supporting the activity characteristic. In such a model, locality means that the agents representing the ants are associated with one place agent at a time, and can only move from one place agent to an adjacent one in the environmental network.



Furthermore, it is not always true to describe the ants simply as agents. For example, a physical ant may be represented by a software ant in an automated system. We call an agent that represents an entity in the real world, an *avatar*. The Hindu meaning of this term is an incarnation of a deity; hence, an embodiment or manifestation of an idea or greater reality. Our avatars are the embodiment in an agent-based system of a greater reality, namely, the physical entities that they represent (Fig. 12). This use of the term is consistent with the use in virtual reality systems, in which an avatar is a computer representation of an external physical entity. [Parunak, 2001b].

Another example of avatars in an environment is mobile agents in a communications network. The network nodes with their interconnectivity form the environment, and the mobile agents are the avatars.



7. Postlogue

By 2015, the social computing is expected by some to morph into ecological or symbiotic computing. John Seely Brown, chief scientist of Xerox suggests that structural matter (atoms) and computing (bits) will become inseparable.

Zillions of sensors, effectors and logical elements (made of organic and inorganic materials) will be interconnected via wireless, peer-to-peer technologies, producing smart, malleable stuff used to build smart appliances, buildings, roads and more. It is during this era that computers disappear. In their place, nearly every physical artifact harbors some computationally based brainpower that helps it know where it was, what was near it, when it was moved and so on. In a way, the inorganic world took on organic properties, using computing to transparently modulate responses to the environment. [Brown, 2001]

8. Bibliography

Brown, John Seely, "Where Have All the Computers Gone?," *Technology Review*, 104:1, February 2001, 2001, pp. 86-87.

Buderi, Robert (2001) "Computing Goes Everywhere," 104:1, *Technology Review*, pp. 53-59.

Castelfranchi, Cristiano (1998) "Modeling Social Action for AI Agents," 103:1-2, *Artificial Intelligence*, pp. 157-182.

Castelfranchi, Cristiano (2000) "Engineering Social Order," *Engineering Societies in the Agent World*, Springer, pp. 1-18.

FIPA (1998) *Foundation for Intelligent Physical Agents FIPA98 Agent Management Specification*, Geneva, Switzerland, Oct. 1998. (<http://www.fipa.org>)

Genilloud, Guy (1997) "Flexible Translation for Integrating CORBA and OSI," *Proceedings of Distributed Object Computing in Telecommunications (DOCT 97)*, Object World, Frankfurt, Germany.

ISO (1994) ISO/IEC and ITU-T, "Information Processing Systems, OSI Reference Model: The Basic Model," Standard 7498-1, Recommendation X.200, 1994.

Odell, James, H. Van Dyke Parunak, and Bernhard Bauer (2000a) "Representing Agent Interaction Protocols in UML," *First International Workshop on Agent-Oriented Software Engineering 2000 (held at the 22nd International Conference on Software Engineering (ISCE))*, Paolo Ciancarini and Michael Wooldridge, eds., pp. 58-65.

Odell, James, H. Van Dyke Parunak, and Bernhard Bauer (2000b) "Extending UML for Agents," *Proc. of the Agent-Oriented Information Systems Workshop at the 17th National Conference on Artificial Intelligence*, Gerd Wagner, Yves Lesperance, and Eric Yu, eds., Austin, TX, pp. 3-17.

Parunak, H. Van Dyke (1997) "'Go to the Ant': Engineering Principles from Natural Agent Systems," *Annals of Operations Research*, volume 75, pp. 69-101.

Parunak, H. Van Dyke and James Odell (2001a) "Representing Social Structure using UML," *Proc. of the Agent-Oriented Software Engineering Workshop, Agents 2001 conference*, Paolo Ciancarini, Michael Wooldridge, and Gerhard Weiss eds., Agents 2001 conference, Montreal, Canada, Springer.

Parunak, H. V. D., S. A. Brueckner, J. Sauter, and J. Posdamer [2001b] Mechanisms and Military Applications for Synthetic Pheromones. In *Proceedings of Workshop on Autonomy Oriented Computation*, 2001. (good example of places, avatars, ...)

Pfeifer, Rolf and Christian Scheier (1999) *Understanding Intelligence*, MIT Press, Cambridge, MA.

Russell, Stuart and Peter Norvig (1995) *Artificial Intelligence: A Modern Approach*, Prentice-Hall, NJ.

SRI (2000) http://www.icann.org/tlds/geo1/ModCE/E_TLD_POLICIES_FIN.htm

Sims, Karl (1994a) "Evolving Virtual Creatures," *Siggraph Proceedings*, July 1994, pp. 15-22.

Sims, Karl (1994b) "Evolving 3D Morphology and Behavior by Competition," *Artificial Life I. Proceedings*, R. A. Brooks and P. Maes, eds., MIT Press, Cambridge, MA, pp. 28-39.

Want, Roy and Bill Schilit (2001), "Expanding the Horizons of Location-Aware Computing," *IEEE Computer*, August 2001, pp. 31-34.

Weiss, Gerhard, ed. (1999) *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, Cambridge, MA.

Wilson, Edward O. (1975) *Sociobiology*, (abridged edition), Belknap Press, Cambridge, MA.