Part III : Cognition
[Belief Agents]

Agents, actions, mental organisms and anigrafs are often confused because the definitions are not precise. A reductionist distinction is that an action is initiated by an agent, whereas a mental organism coordinates a sequence of actions carried out by several agents. An anigraf shows the relations among the goal states for agents or mental organisms, but sometimes also identifies a group of agents or organisms. To the external observer, a creature's actions are used to infer intentions, or often a state-of-mind. Such states-of-mind represent the beliefs and goal relationships of the creature and are reflected in the structure of the anigraf. These relations may be context-sensitive, in which case the anigraf form may change even if the goal states remain the same.

In the natural world, creatures interact with one another. The nature and success of these engagements depends upon the internal models of the interacting participants. Mental organisms that hold beliefs about the anigraf forms of other creatures will extend the anigraf abstraction. The anigraf now encompasses more than a physically encapsulated system. There will be additional mental organisms attempting to control actions of others as well as oneself. Actions provide information about an intentional stance. Interpreting actions, however, requires recovering the belief structures of other anigrafs.

1 Aug 07

Anigraf4: Dancers
mating games

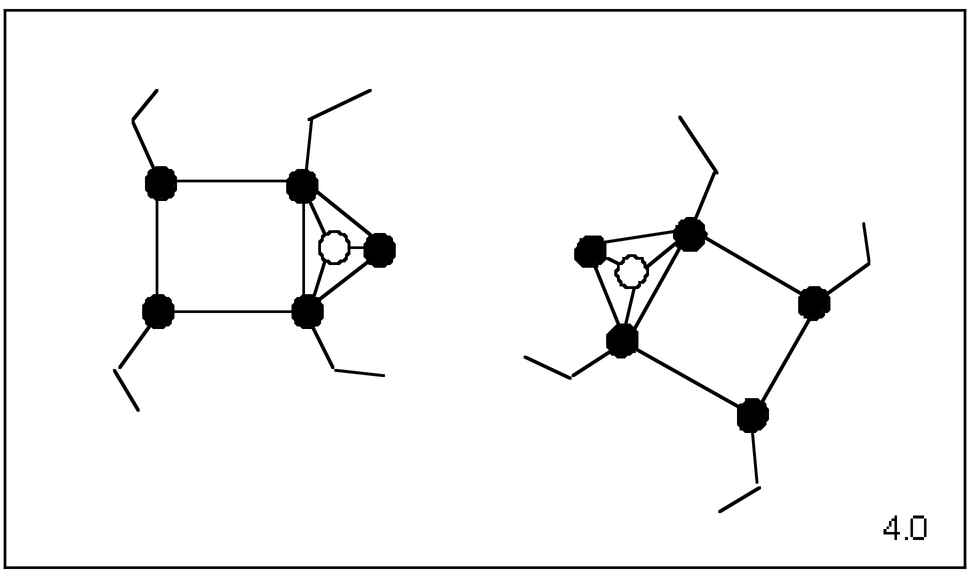

Fig. 4.0. Two legged anigrafs that are poised to interpret the intentions of the other.

1 Aug 07

4.0 Two to Tango

Goals such as "fight", "flee", "evade", "approach", or "be still" have a definite cognitive flavor. Lying beneath these goals are the gaits we characterized as run, walk, turn, halt, etc. At each of the two levels of description, we have different, but related anigrafs. Let us attribute the intentional stance to the more cognitive, mental organisms, and the lower level, gait control to agents. The external observer sees only the gaits, and from these actions infers a creature's intentions and goals. For example, the direction and pace of a walk might suggest an exploration if forward, or timidity if backward. Such movements thus provide elements for a visual language. But to understand this language, there must be a correspondence between anigraf models. The complexity of the language increases as the actions of other parts of the body come into play, such as head nodding or the waving of limbs. Fortunately, in the world familiar to us, a set of rather elementary movements is prevalent across many animal forms (Davis et al, 2000). They include bobbing, circling, or swaying of one's body, limbs or head. When two creatures of the same species coordinate such patterns of movements, we have a dance.

Definition: *A* **Dance** *is a sequence of movements, such as gaits, executed in coordination by two physically distinct entities.*

The challenge is to understand how the dance can be coordinated at the level of the anigraf, where two sets of mental organisms must recognize the next (or present) goal state of the other. A simple courtship sequence provides an example.

4.1 Beyond Reflexes

It is almost trivial to build automata that execute dances. A pair of "lego robots" can perform a "cha-cha" using distance-sensitive IR sensors that control the proper forward and backward movements needed. However this solution is a simple feedback controller. At issue, therefore, is what makes our latest Anigraf4 different from robotic systems, whose behaviors are no more than what one might expect from a pair of Braitenberg's vehicles. How can a social awareness common to both partners be put into place? In effect, we will need to merge two separate packages of mental organisms and their corresponding anigrafs into one integrated, social decision-making system.

To create a dance with a social awareness between the partners, a minimum of three conditions must be met:

> (i) Each partner must have a repertoire of movements, or patterns of movement in common, with similar anigraf models relating these movements.
>
> (ii) There must be a communication channel between corresponding agents or mental organisms that comprise the two anigrafs. By corresponding agents, we mean agents who have identical preference rankings for the various possible movements. This identity means that each corresponding agent or mental organism understands the intentional stance of the other (but not necessarily of all other mental organisms.)
>
> (iii) When each anigraf conducts a tally, the tally must include the votes of the mental organisms (or agents) in the other anigraf. This condition insures that the decisions include both parties, and hence include an element of social awareness.

The first requirement is easily filled if each partner can execute a common set of movements. These movements then become the goal states for a corresponding set of mental organisms residing in the two identical anigrafs. For a dance, these goal states should reflect behaviors such as flight, fight, avoid, approach, etc. and not as a sequence of specific limb movements. For example, one goal might be "to kiss" the partner; another might be "to back away" (because of shyness or fear.) For very sophisticated anigraf dancers, the choreography of such movements has the potential of telling a story (Heider and Simmel,1946).

To fullfill the second, communication requirement, corresponding agents, each associated with their own anigraf, must be linked with some kind of channel. This linkage could be visual, acoustic, tactile --whatever-- as long as the channel can convey the strength of preferences of the linked agents. This stipulation is a bit more complex than the simple sensory channels found in most robots. Here, the communications must convey preferences for the actions of the system -- i.e. what dance steps the two sets of mental organisms should be voting upon. Agents engaged in the linkage can then pass on their assessment to higher-level mental organisms, who make the final decision.

The third requirement for an animate dance is the most crucial. The final choice of dance step must apply to both anigrafs as one integrated social system. Hence the tally for the preferences for actions should include all agents or organisms, combining votes from both parties. But here lies the problem: there is no social aggregation process. Specifically, there is no tally machine common to both. Rather, each set of agents is forced to make its own tally. At some abstract level, then, our dancers would be no different from the robots. Neither party would be performing a social activity that involved a decision made by aggregating the preference states for the entire group of mental organisms engaged in the dance. If anigrafs are to support social dancing, then each tally must include the preferences of the partner.

4.2 Proxy Votes

For a mental organism or agent in one anigraf to be a part of another's tally, the partner must acknowledge the corresponding agent's opinion as being as valid as that of its own internal agents. To solve this problem, create a new type of agent called "a proxy." Proxies are agents who reside in one anigraf, but vote the wishes of another. A proxy agent of type P will be designated using italic notation P'_{M} or P'_{W}, with subscripts added when necessary to identify the anigraf, such as "M" and "W". Obviously each proxy must share a communication channel with their counterparts to enable them to read the other's preferences. So, for example, if one anigraf "W" has reached the decision "to kiss", then this intent would be read by the "kiss" proxy agent K'_{M} associated with anigraf "M". Such intents might be indicated by a pouted mouth or by an advanced head position. Similarly, a head lift or turn might indicate a decision to move apart or to create a circular movement, and again read by another proxy agent T'_{M}. Each proxy agent K'_{W} or K'_{M} would give anigrafs the potential to know the other's intent for the next movement. To complete this potential, however, we may require that the *P'*-agents be linked in the same way as the action agents A in the anigraf proper. Then the set of each *P' s* will have similar intrinsic knowledge about just what these goal states signify. Note that this constraint is no different than that imposed for our most primitive anigrafs: namely edges in the graphical representation indicate that an agent "understands" goal states for the system that are different, but similar to, its own and then includes these options when setting preference orderings.

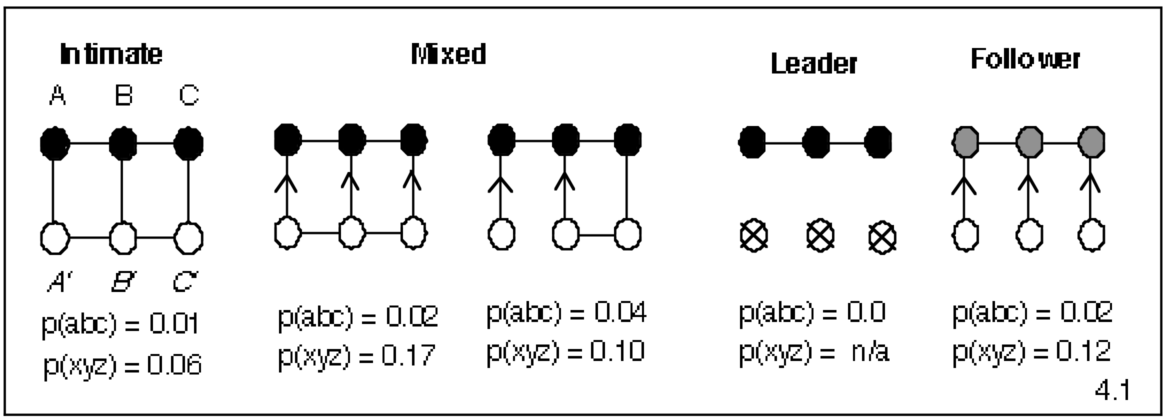

Fig. 4.1. Examples of relations between proxy agents (open circles) and action agents (solid circles) in an anigraf. Probabilities for abc cycle are indicated, as well as for a cycle between any three nodes, including proxy nodes.

The proxy - to - agent relation can have a major impact on an anigraf tally. To illustrate, five simple networks for one anigraf are shown in Fig 4.1. Action agents are filled circles; the analogous proxies who report the states of a second anigraf are open circles. Here the subscripts are omitted, because we explore only the relation between proxies {*A', B',C'*….} and their corresponding action agents {A, B, C…} in the same anigraf. Consider first the two extreme cases at the right, where the network design dictates whether a partner will be a leader or follower. For the leader, proxy weights have no influence at all. Yet for the follower, the proxy weights will dominate the weight total for each of the action agents. (See related consequences in inset box.)

Another extreme is when both anigrafs share intimate, tight bonds, meeting each other halfway, with bilateral edges between proxies and agents. The depiction is illustrated at the far left of Fig. 4.1. Between these extremes are various combinations, two of which are shown as networks with a mixture of bi-directional and directed edges.

Hypnosis and Hallucinations: two byproducts

Two fascinating situations arise when an anigraf becomes subservient to the proxy demands. In both of these cases, the proxy communication channels are active with strong inputs, but the action-agents themselves would be very weak. Let one anigraf's desires dominate the proxy channels of the other. Then that anigraf is completely controlling the behavior of the other. In a second version, we could imagine random external inputs controlling the proxy values, but with no second, dominant anigraf present. The situation is now like a dream, or hallucination, whereas the first resembles an hypnosis.

Could there be an evolutionary advantage to losing some control over one's active agents? For example, what if the current tallies lead to chaotic, maladaptive behaviors? Or what if there is some kind of frustrated decision making, incompatible with the physical world? Reliquishing control by active internal agents might then be beneficial. Viewed in this light, proxies could be used to produce a calming effect on a creature's behaviors. Examples of such beneficial proxy control might include grooming, stroking one's skin, rocking motions, soothing music, etc. Under the right circumstances, suitably tailored proxy agents could easily dominate an anigraf's tally. Could dreams, hallucinations, and hypnotic states have a set of (proxy) agent-types in common?

The probability of cycles for various proxy-agent networks provides further insight into the power of proxies. In Fig 4.1, the probability of a top cycle among the agents {A, B, C} is p(abc), whereas the probability of any top cycle in the network p(xyz). Not surprising, complex networks with directional proxy edges tend to have more top cycles. When the entire network is bi-directional, as shown in the ladder graph at the left, cycle probability is low. Now the state of the action agents {A,B,C} contributes directly to the proxy votes {*A'*, *B* ,' *C'* } when a tally is taken. This is a curious twist: if an anigraf's own beliefs influence how another anigraf's intentions are "read", then non-cooperative behaviors may result even for otherwise stable networks. This seems inappropriate. Hence, regardless of the fact that top cycles may increase, the more plausible proxy-agent networks should have directional edges so the strength of proxies will influence the weights of action agents, but not vice versa.

Regardless of the stability and instability of outcomes when proxies participate in a tally, these new agents add the ingredient of a social awareness held in common between two partner anigrafs. In the ideal case, when the partners are twins with the same anigraf model, together with perfect proxy readings, then there is a strong "meeting of the minds." Each anigraf has been augmented in an equivalent manner such that when the Condorcet tally is conducted *in parallel, within each anigraf* , then the outcomes will reflect the intentions of both anigrafs. A shared social consciousness emerges.

4.3 A Duet

Figure 4.2 shows a more detailed abstraction of two Anigrafs who might participate in a duet. The solid circles represent the internal agents who effect actions, and the open circles are the proxies. The actions are to move forward "to kiss" (K), to stop the kissing (S), to move apart (B), and to halt the backward movement (H). These actions are related as illustrated in the figure, which depicts the underlying "mental model". This model is equivalent to a 4-chain of agents, plus three associated proxies that complete a ring centered on a merged proxy *S'/H'*. Cyclic outcomes supporting a coordinated dance are now easy to arrange.

As an example, consider a cyclic "cha-cha" that has four action states: K, S, H, B as in Fig 4.2. With directed proxy input, a three cycle will occur by chance about 20% of the time, and the rarer four cycle occurs for about 5% of random weight choices. Whether or not all or none of the proxy agents are

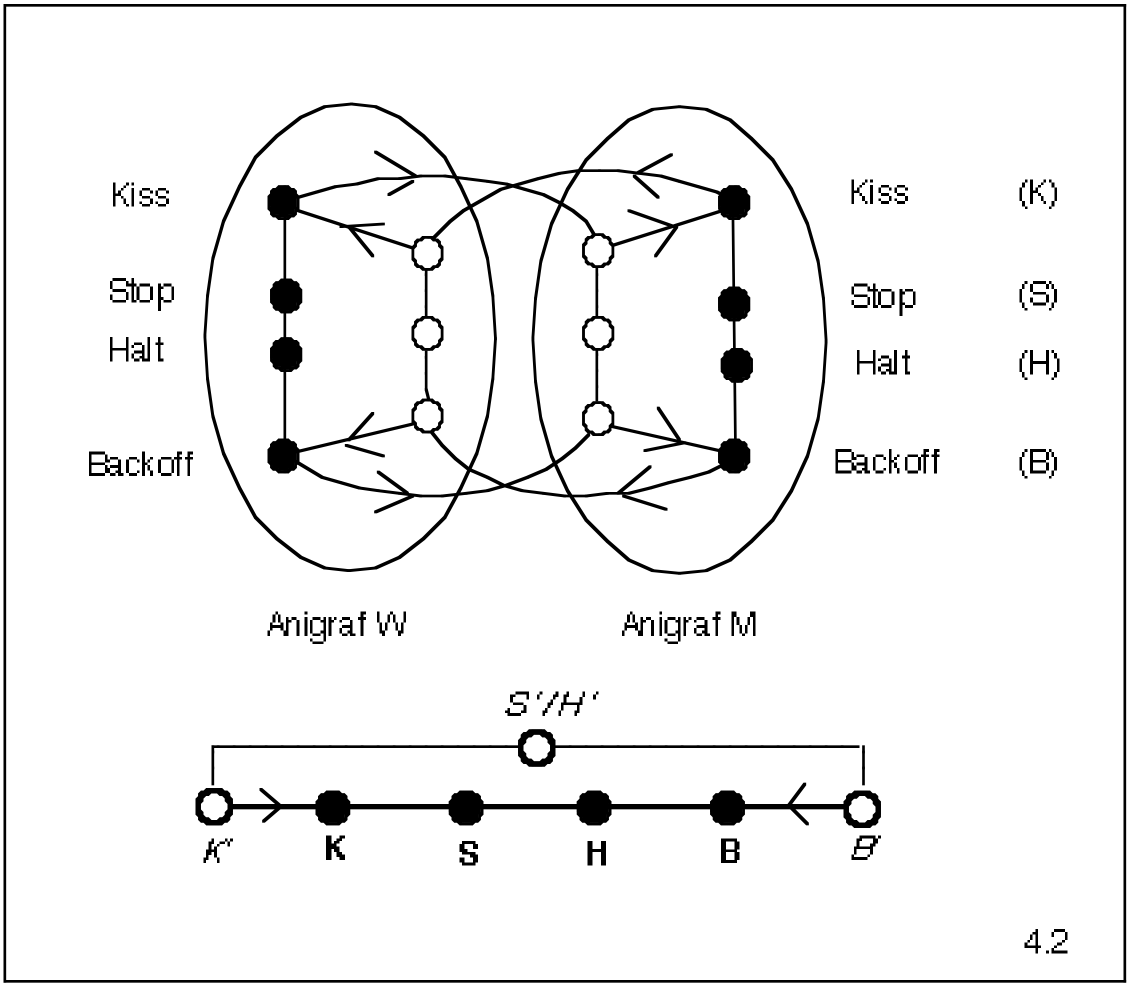

Fig 4.2 Two anigrafs with identical sets of proxies and agents. Note that all agents do not have unique proxies in this example. The lower graph clarifies the structure for each anigraf.

connected to each other makes little difference. These percentages are high enough to allow a simple dance to be learned, even with trial and error searching for appropriate weights on nodes. Fig. 4.3 shows one region of top-cycle activity that exceeds 10% when weights of both agents and proxies are chosen from a uniform distribution. Plate 3 shows another slice, revealing more clearly the detailed structure of part of this space of outcomes.

One might now ask again how this “cha-cha” differs from one enacted by lego robots or vehicles? First, the complete “cha-cha” requires a K,S,B,H 4-cycle in a graph that need be only partially directed. Because there is bilateral communication between action agents, with shared information included in each tally, the outcome a social process. Although the proxies might be regarded as robotic-like sensor inputs, their effect is quite complex and not modeled by classical controllers. The state transitions arising from the pair-wise tally are

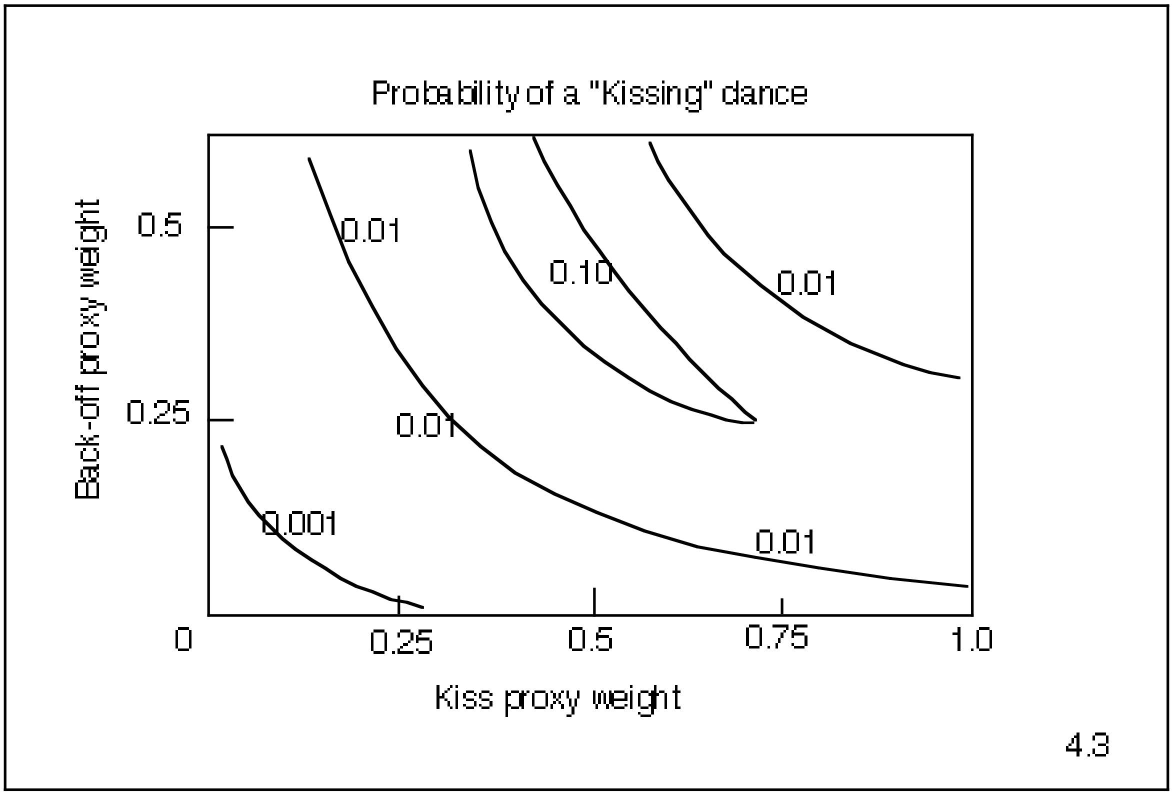

Fig. 4.3. A slice through the six-dimensional space showing the probability of top cycles for various proxy weights.

very difficult to predict. Outcomes are not simple additions of an agent's voting power and the proxy weights. For example, our paired anigrafs *could* engage in a cyclic cha-cha even if agents *K'* or *B'* but not both had zero voting power! Alternately, if both votes for both *K'* and *B'* are simultaneously near their maximum weight, then there will be no dance cycle. This makes intuitive sense: mental organisms can not collectively decide to engage in kissing and breaking up simultaneously. The cyclic dance, then, succeeds only with some measured input by both sets of proxies. Anigraf social behavior is quite non-linear.

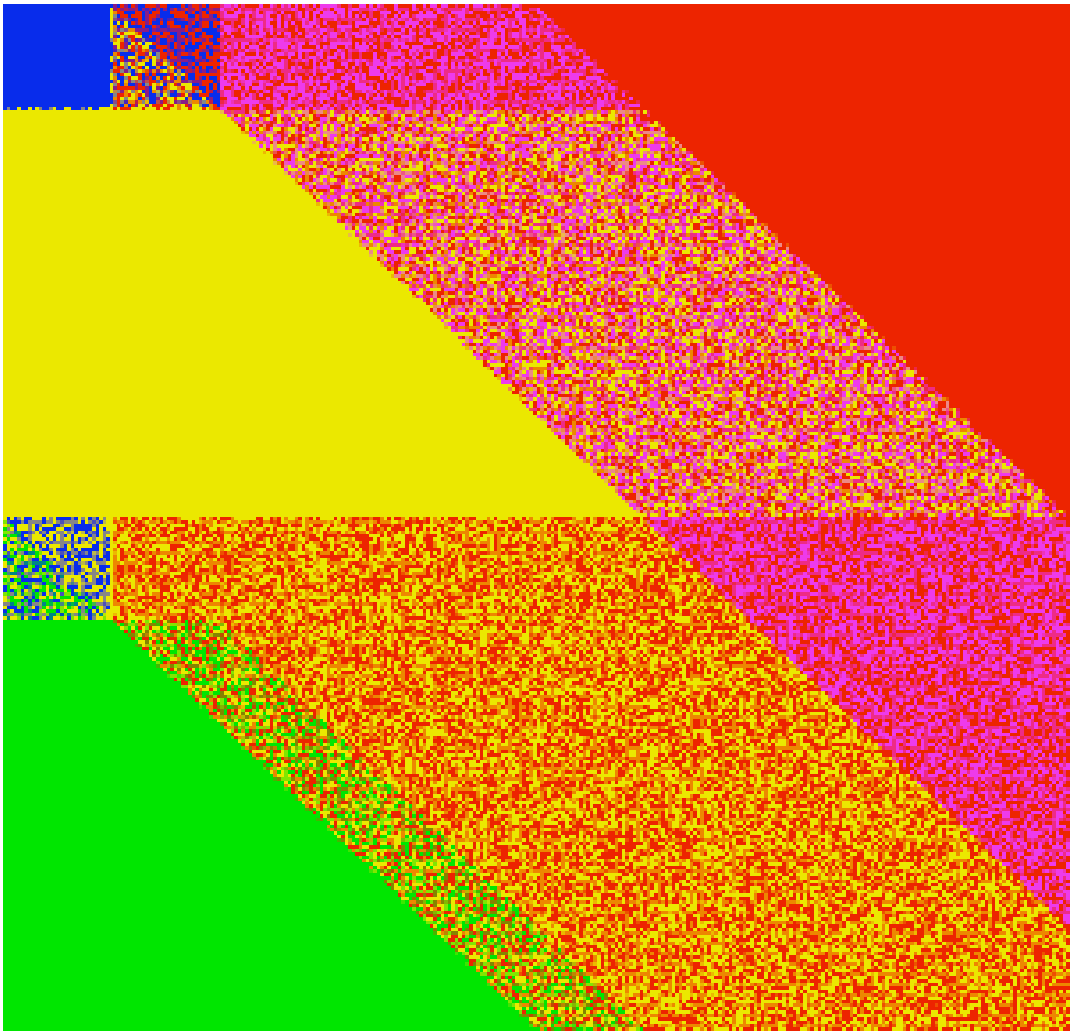

Plate 3: Phase Plot for a Kissing Dance of four steps (equivalent to a directed chain with two directed proxy agents at end of 4-chain.) Textured areas are top cycles. Weights are (4, x = 0,10, 5, y = 0 –10, 0, 6.) Note regions of 4-cycles.

4.4 Social Twins

Most creatures use many different channels and modalities to communicate. Bees, for example, have at least three different channels to indicate food location, quality, and quantity: body motions (two different patterns that include body contact), wing beats, and sounds. Ants have an array of chemical pheromones. The number of different signs used by animals is almost countless (Darwin, 1859). Adding more specialized communication channels allows the coding to be greatly simplified, in the limit reducing to sending a single yes/no bit. But this takes place at the expense of adding more proxy agents. The great benefit, however, is that the aggregation process becomes more social: each anigraf now has greater participation in the partner's tally. Schematically, two interacting anigrafs of the same species might be depicted as in Fig 4.4. Here, for clarity, only three proxy communication channels are indicated by dashed lines.

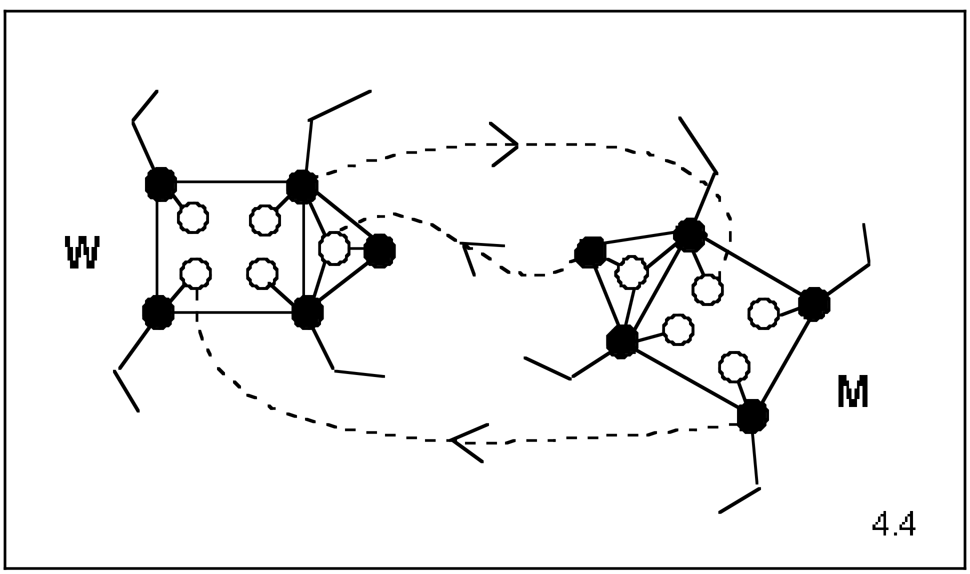

Fig. 4.4. Depiction of the interaction between twin anigrafs. The solid nodes are action agents, the open nodes are proxies for corresponding action agents in the twin, to which they have virtual links (dashed lines.)

(Line segments without nodes are included to indicate limbs controlled by lower-tier agents.) Like all other anigrafs, **M** and **W** each have their own separate tally machine. But each tally will yield the appropriate social outcome for the actions of the pair. At least this will be true when **M** and **W** are identical twins who are initialized with the same weights on identical action agents. However, even with slight differences in anigraf forms and proxy reliability, anigrafs now have a sense of understanding the intent of others, allowing them to engage in a variety of group activities. These are first steps toward social communication.

Anigraf5: Workers
rolling stones

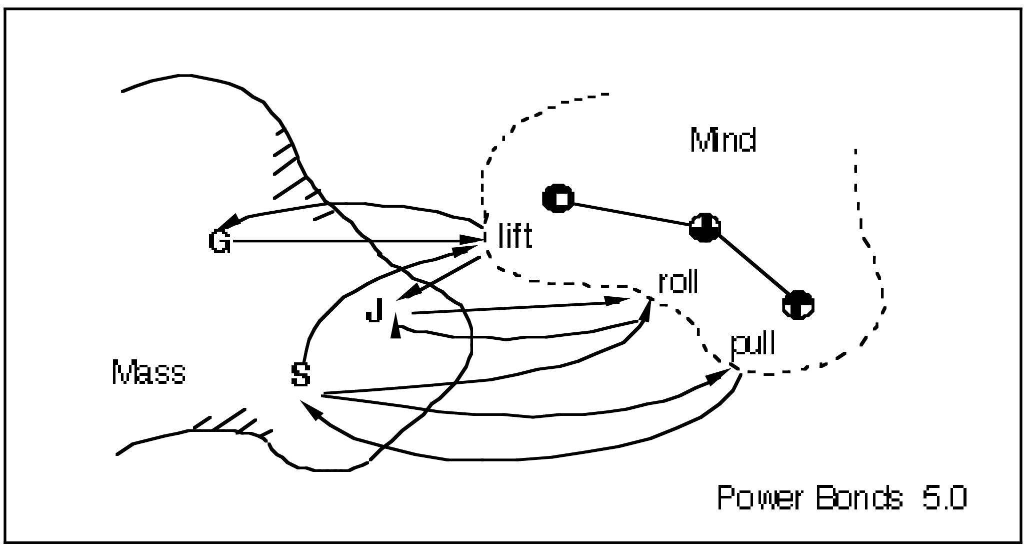

Fig. 5.0 Schematic showing an anigraf (left) with three goals: lift, roll, pull. The mass on right has three types of opposing forces: gravity (G), sticktion (S) and inertia (J).

1 Aug 07

5.0 Embodied Anigrafs

We have danced between anigraf abstractions and hints of creature constructions that include body or limb motions. When the objective is to coordinate the mental goal states of two anigrafs, as in a Dance, the actual form of the physical world is not a major issue. But when an anigraf's task is to manipulate external physical objects, then the mental organisms that comprise the anigraf must have a related set of intentional acts that make sense in terms of the world they wish to manipulate. The physical world most closely associated with the anigraf is its own body. Consequently, the nature of an anigraf's own embodiment becomes the most important factor in building a model for its physical world.

Ideally, if structure and function were correlated, similarity relationships between mental organisms and the parts of the body they control would be expected. In this case, different creature designs will lead to the development of different mental representations, or anigraf forms, for actions and events in the world. Although biological forms and mental representations may differ, these forms are not arbitrary. Biological structures tend to have sets of properties that are highly correlated (Thompson, 1968; McMahon, 1977.) In other words, observing one property is highly predictive of another. These are so-called Natural Modes, which provide a probabilistic view of Nature's essence (see Appendix 5.) This fact is important in the construction of anigraf models that need to make explicit which properties (and behaviors) in the world are closely associated. Predictive anigraf forms thus begin with a model of relations among their own sensory and motor abilities (of which it is aware.) From these, the dimensions of events and objects emerge. Hence the dimensions of physics are not Nature's choices, but are subjective constructs of mental models. Stated more strongly, the choices of what we call the "physical dimensions" depend on measurement abilities, models, and context. As Max Planck aptly states: "to inquire into the "real" dimension of a quantity has no more meaning than to inquire into the "real " name of an object" (Langhaar, 1951).

The dimensions of the external world of anigrafs are thus grounded in their embodiments. If a creature moves, then it is likely to have a sense of time, space, and/or velocity. The time-piece could be a gait cycle, a heart beat, or even a tally rate for one or more sets of mental organisms. If seasonal or daily changes are sensed, then yet another clock can be inferred. Speed might be measured by gait cycle rate, and distance could be derived as units of time, or simply a unit of body size if a smaller scale is needed. If a creature has sensors for heat loss, then associated proxies could provide the basis for a temperature scale. Thus we have the beginning of some dimensional constructs for a physics, namely time (T), length (L), and temperature (Δ°). Our most important unit, however, will be the analog of effort, namely the product (FT) of force and time. (Mass can be recovered if necessary using Newton's second law: e.g. F has dimensions MLT^{-2}.)

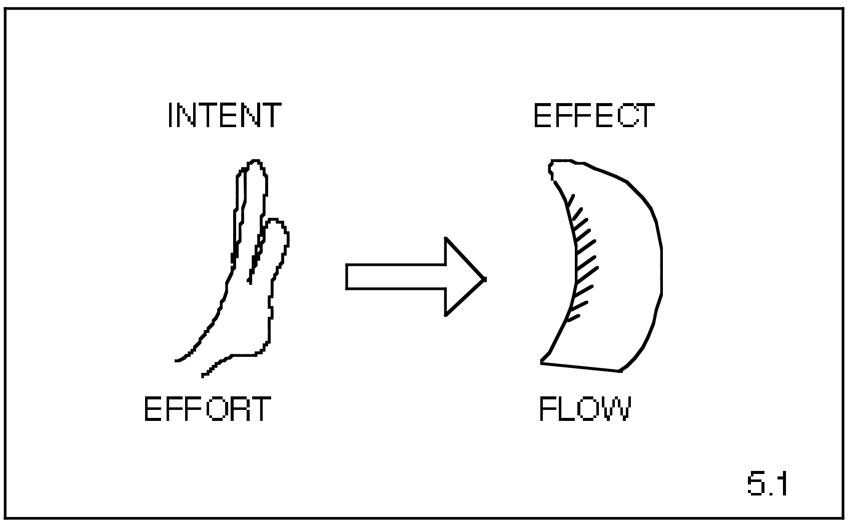

A crude depiction of the relation between intent and physical force (or effort.) A critical point is that the flow of power goes two ways. Not only does the hand push the object, but the object is also pushing back on the hand. The difference being that on the one side the force is an "intent", whereas on the other, the "force" is associated with a mass.
The intent (internal force) over time is *directly* analogous to power (FT), which, when applied to an object will effect some movement or "flow" (MLT^{-1}). The extent of "flow" will be dependent upon the resistance of the object (here M), and the resultant velocity of object motion (LT^{-1}). This simple "cause-effect" relationship is ubiquitous though out the natural world (e.g. E = IR, Q=M£, P=Fv) and has been elegantly captured by Paynter (1961) using "power-bond graphs."

Effort is experienced when interacting with external objects (including one's own limbs and body). These experiences lead to a compelling intuition captured by two of Newton's three laws: (1) objects at rest or in motion maintain those states unless acted on by some force, and (3) when one body exerts a force upon another, that body receives an opposite (and equal) force. Already, however, we have seen a psychological analog: To change opinions and to reach agreement, anigrafs require an expression of the strength of *intentions* and desires to overcome those of unwilling partners. Specifically, changes must be induced in the weights or form of the other anigraf. The necessary "force" invoked is no different from the "force" of an intention to make a limb movement, where the exertion has a physical effect. Efforts (FT) and intentions are intimately coupled. In both the physical and psychological domain, these intentional levels are monitored as part of the tally process. If the strength of intent is sufficient, then the desired action will take place. This relation between intentional effort and consequent action, or flow, is captured elegantly by Paynter's Power Bond graphs (see Fig 5.1 and inset.)

5.1 Inanimate Intentions

When two anigrafs coordinate their activities, the strengths of the intentions of the resident mental organisms are being continuously monitored by proxy agents, with the proxy's assessment of weights affecting outcomes. If one anigraf seeks to dominate another, then additional effort must be expended to obtain submission. Manipulating physical objects is no different. In each case, the common currency is the intentional effort needed to bring about a consequence. From the anigraf's perspective, any physical object "pushing back" is no different from another anigraf with intentions.

Imagine an inanimate boulder, which the anigraf wishes to "dance with" -- i.e. to move. Referring to Fig. 5.2, the anigraf's initial model will be the same as that for another anigraf: there are action agents (lift, roll, pull..) and their analogous proxies (*L', R', P'*...). Already the anigraf must have some notion of mass (M) if it can lift objects (it's leg for example.) Hence, although the gravitational constant G is never directly sensed, it is inferred as a "constant of the context", as also are inertial (J) and sticktion (S) parameters. These unseen constants become embodied in the anigraf's correlations between its own intentional powers (for lift, roll, pull..) and resultant actions. Nature's Modal designs insure that such correlations lead to useful and stable embodiments of these kinds of correlation parameters (again, see Appendix 5.) However, all that the anigraf need know initially is that the *in*anigraf object also will try to "roll", or "push", or "pull-back", and that a sense of the strength of these actions will be reported by the proxies. This is no different from the mental organisms of one anigraf learning to coordinate a dance with another.

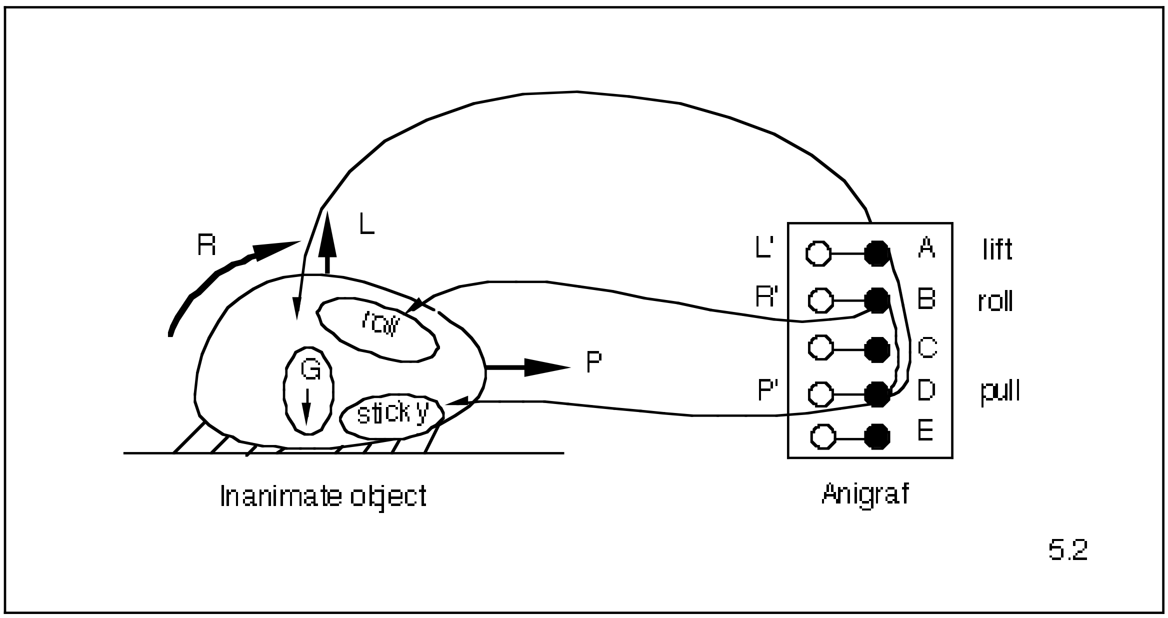

Fig. 5.2. A simple anigraf with three of its five states in play: lift, roll, and pull. These states comprise the creature's model for the inanimate object, similarly assumed to have only three agents with comparable, but opposing goals. The anigraf senses the object's intent via three proxies, *L', R', P'*.

5.2 Analogical Physics

The naive anigraf's first model of inanimate entities is that this type of object is composed of recalcitrant agents. The proxies are reporting intentional levels of this *in*anigraf-object, but after a tally is taken, the expected behavior is not observed. Unfortunately, however, a boulder does not have the same internal model of the relationships between "pull, roll, lift..." as that of the anigraf, and has no interest in any "social" behavior. Hence there can be no "cooperativity" following a tally. The boulder simply "does its own thing" -- to seek stability under gravity and inertial forces.

To give one specific example, consider again the boulder depicted in Fig. 5.2. The goal is to roll the object. The first task of the anigraf will be to evaluate the "intentional" forces of all the component masses of the object. Inertial mass, or more properly, resistance to Paynter's "flow", can be estimated either by size (if experienced in evaluating densities), or by actually probing the object. This proxy reading is then compared with the intentional strength of the anigraf's own action agents. In this simple case, the vote must have a majority favoring the action agents. A more sophisticated or determined anigraf might probe the boulder in several places, to see if the intentional "forces" of the masses are the same in all directions. Perhaps there is a weakness in one of these directions. If the boulder is wedged, however, more than inertial moments alone will come into play. If lifting and rolling are required simultaneously, then the summation of proxy information will be comprised of two quite different kinds of forces -- yet these will be indistinguishable to the anigraf, who only needs to evaluate the common currency of "intent" -- as if the boulder simply encapsulated another obstinate anigraf.

To predict the outcomes correctly, the anigraf must thus recognize that the context has changed. Once noted, one simple trick would be to take a "brute force" approach, thus requiring only one physical state to be evaluated (e.g. to "lift"requires that only "mG" to be overcome.) However, if only two or at most three physical states are coupled (e.g. lift and roll), then something akin to a vector sum must be calculated. The nearest social procedure is the revised Borda method presented in the Preliminaries (see Table 0.1). Fortunately, it is easy to show that this (revised) Borda Count and the Condorcet outcome will be the same for any physical system with three or fewer independent states such as illustrated in Fig. 5.2. Another trick might be to seek the Condorcet (or Borda) loser. Certainly this state would be the easiest to overcome, but would require a rather dramatic alteration in the customary tally.

5.3 Machines

Levers and wedges are two of some of the first machines invented by Man. Each amplifies the application of forces, reducing effort. A nut is cracked near the hinge of a jaw; a two segment limb can lift a heavier object when the elbow is at a right angle, rather than at 180 degrees when the limb is completely extended. More power is achieved if the elbow is used as a fulcrum, being supported by a ground plane. But to fully grasp the concept of a lever as a rigid rod positioned on a fulcrum came only in recent millennia (roughly 5000 BC.) Perhaps the conceptual origin came from playing with sticks, such as noting that advantage of throwing a stick from one end, rather than from the middle. Such explorations and conceptualizations transcend our primitive anigrafs. They are amplifications of primitive technologies – both physical and intellectual. A new world for anigrafs become possible. This will be part of Metagrafs presented in Part IV.

5.4 Summary

Our anigrafs begin their education of interactions with inanimate objects by using a model which is a copy of the model used to understand intentions of other members of their species. This model is thus identical to its own (directed) internal model for understanding how effort is related to the movement of itself and its body parts. However, there are several important differences:

(i) inanimate goals are typically “simple” (e.g. stability.)

(ii) inanimate agents are totally honest. Hence intentions can be read accurately up to the ability of the interface.

(iii) there is never any conscious cooperation.

(iv) all inanimate models will be directed graphs, typified by Bond graphs.

(v) the Condorcet tally machine can still be used with success, but for no more than three inanimate goal states.

We can expect the adaptive anigraf to recognize the inanimate object's failure to engage in any cooperative activity, and to use this difference to identify the context. But how will the anigraf evolve new, more powerful models for any context? A first step is to understand how experiences with new and unfamiliar states can be incorporated into current models.

Anigraf6: Explorers
new worlds

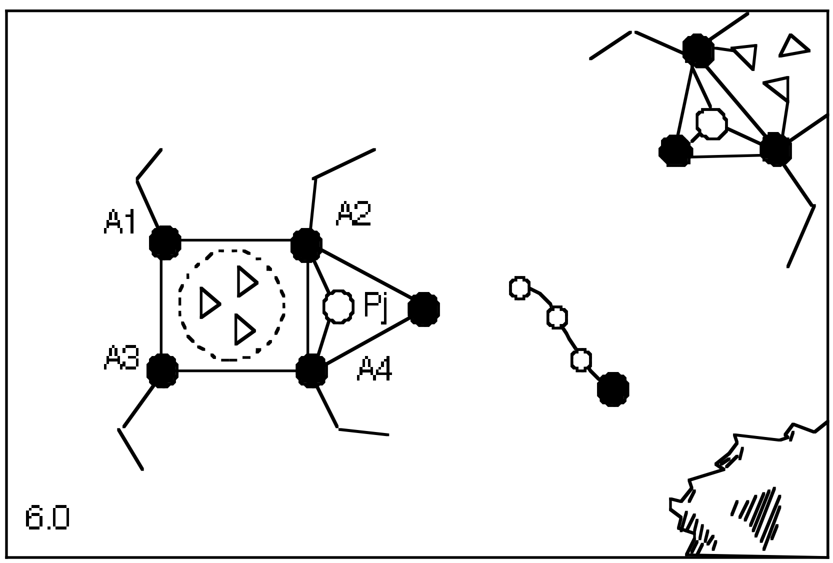

Fig. 6.0. Two anigafs inspecting a third, worm-like creature. Action nodes are solid circles, proxies open circles. Triangles represent free agents, some of whom have become linked to action agents in one of the anigrafs.

3 Aug 07

6.0 Expanding Horizons

As anigrafs become engaged in manipulating their worlds, new relationships among events will be discovered. Some of these involve inanimate objects and actions; others may follow from interactions with new species of creatures, perhaps with different embodiments and hence with different internal models. These new relationships will require new, or revised models that are appropriate for the context.

Anigraf networks can change or evolve in two obvious ways: edges can be added or deleted; or, nodes can be created or destroyed. Both come at a cost. If either edges or nodes are revised, then preference orderings appropriate for the new context may be misleading in the earlier context. Even the addition of a single new node creates problems. Actions associated with this new node become second choice preference for all the original nodes to which that new node is connected. Preference orderings are affected. If the degree of connectivity is high, the single node addition can cause major changes in the anigraf network, and hence its model for actions. Nodes and edges cannot be added haphazardly.

In order to avoid disrupting actions of successful anigrafs, one solution would be to duplicate the original anigraf, and then revise relationships appropriate for any new context. If certain anigraf forms were very common (e.g. rings, chains, trees….), it is easy to imagine a reservoir of these forms available for use. To facilitate the development of new models, perhaps each active clique of mental organisms have an associated stack of similar models with unassigned nodes. Then a context revision would simply correspond to a shift to another level in the stack, with an analogical labeling of these new nodes. In effect this was the solution proposed for the worker anigrafs who were faced with using an animate model to describe the behavior of inanimate objects.

A second, more conservative possibility would be to refine the categorical choices of agents, replacing one agent with several alternative agents with subcategories that have become more significant and desirable. For example, whenever sensory inputs or motor outputs are augmented, new agents or mental organisms will be added to represent these new states. In the early stages of development, such changes in an anigraf's form are most likely node additions. This simple inference, although easy to envision, leads to formidable theoretical challenges.

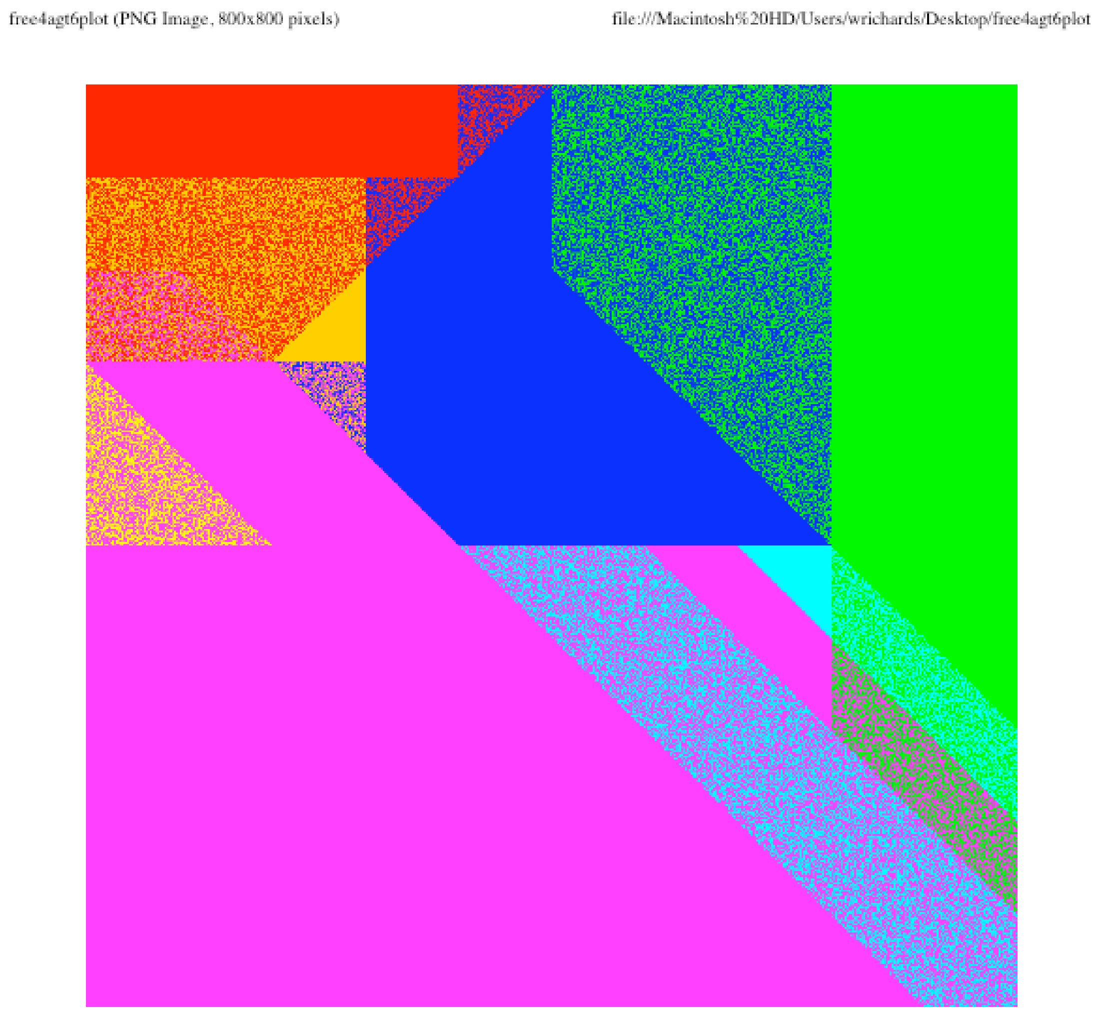

Plate 4: Phase Plot showing Winners for an undirected chain of 6 agents, with 4 free agents of zero weight randomly connected to these agents. Textured areas show topcycles. Some cycles include free agents. Weights are (6, 1, x = 0-10, 2, y = 0-10, 4, 0, 0, 0, 0.)

6.1 Free Agents

To add nodes, the anigraf must possess some "free agents." These are agents whose preferences are initially unassigned, but are capable of acquiring a new goal state and learning its relation to present goals. Hence free agents initially have zero voting power. They will be depicted as triangles. In some cases, several free agents may be grouped together to form a clique, already having tentative (unspecified) relationships to one another. An example is when one anigraf model has a stack of similar forms available, with node assignments to be made later by analogy.

One might assume that if "free agents" have zero clout and no preferences, then they can be connected arbitrarily to action agents without changing behaviors. But this is not true. Consider a single free agent with no behavioral goal and who is bi-directionally connected to a large number of action agents. Then in any Condorcet tally, this free agent, with no goal state, could have a very large weight score resulting from the addition of all its neighbor's weights. This total could exceed that of any other active agent, resulting in a winner that leaves the anigraf in a quandary with no specified goal. The odds for such a stalemate increase as the ratio of free agents to action agents increases. Thus, if anigrafs are to have the potential to acquire new, successful models for actions, then constraints need to be imposed on (a) the size of the population of unlabeled (free) agents with respect to those already labeled, and (b) the form of connectivity between inactive, unlabeled free agents and active agents. (Here we assume some modicum of communication as a precursor to establishing a more permanent relationship.)

Consider a set of six action agents, arbitrarily linked to form a random anigraf *G6*. Now add at random to any of these action agents an increasing number of free agents. As shown in the second row of Table 6.1, if we add 16 free agents with no voting power, twenty percent of the tallies will be won by one of added 16 free agents although they have cast no votes. If the fixed set of action agents has the form of a chain *C6*, then the percent of free agent

Percent of Free Agent Winners					Tbl 6.1	
#Free Agents	0	2	4	8	16	32
Random ***G6***	0	4	9	15	20	25
Chain ***C6***	0	18	28	40	55	70

Table 6.1: Increasing the population of free agents with bi-directional edges increases the probability of free agent winners and hence a "do-nothing" consensus.

winners becomes disastrous even when only four such agents have been added with no voting powers. (See also Plate 4.)

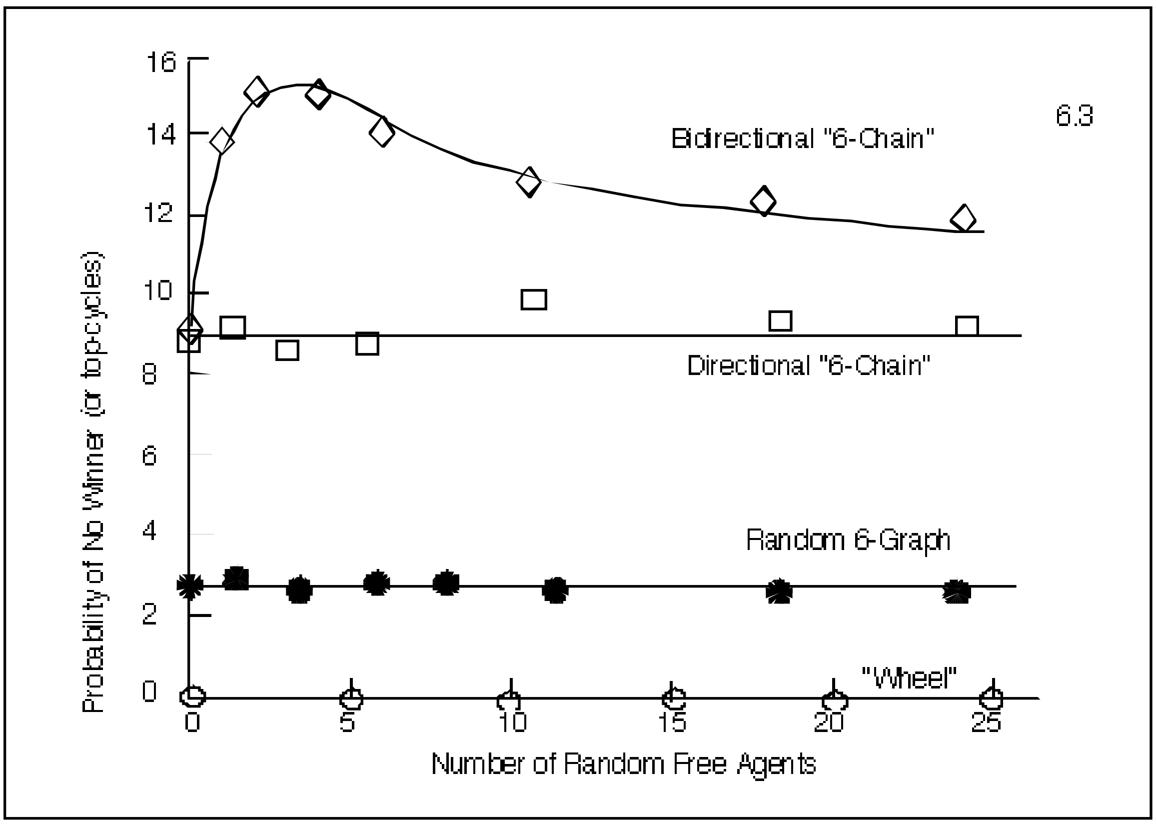

Fig 6.1. The probability of topcycles for four different kinds of 6-node anigrafs, when free agents are added with directed edges. (The random graphs have edge probabilities of one-half.)

A possible remedy is to require directed edges between free agents and action agents. Now an action agent is highly likely to be the Condorcet winner, reducing the possibility of a "do-nothing" consensus. Furthermore, as shown in Fig. 6.1, for random graphs with six action agents, only about 3 % of the tallies will now have top cycles; better yet, covered graphs, such as a "wheel" of six action agents will have guaranteed winners with no top-cycles. A 6-chain however is less resilient to free agent additions, more than doubling the rate of no-winners over random graphs. Note the significant differences that depend on whether free agents are added with bi-directional or directional edges. For the neuroscientist, these consequences may be disturbing: two neural networks that on first appearance appear identical could have quite different behaviors.

6.2 Ideals

What is the ideal number of free agents for any given number of action agents? Three constraints seem obvious: First, every action agent should have associated with it at least one free agent, thus giving this agent some potential for expressing a new relationship. Second, any one free agent should not be associated with a large number of action agents, otherwise any specificity in a new acquisition is diluted and the odds for a free agent winner become objectionable. Third, top-cycles should be minimized or lie within some reasonable bounds (say 10%.) The limited data of Fig. 6.1 suggest a possible satisficing solution of about twice the number of free agents for any set of action agents. Consider first the random 6-agent case. For this set, we need at least six free agents – one for each action agent. However, for random free agent-to-agent assignments, we need to add twelve to give good odds (p ~3/4) that each agent will have at least one free agent associate. Random graphs thus require roughly twice as many free agents to action agents. For special graphs, such as chains, the ratio need be only slightly higher, and for highly covered graphs such as the "wheel", the ideal ratio may be lower. But what about the proxies associated with agents? How should these be augmented? Is there a relation between ratios of free agents, proxies and action agents that will facilitate the model development process?

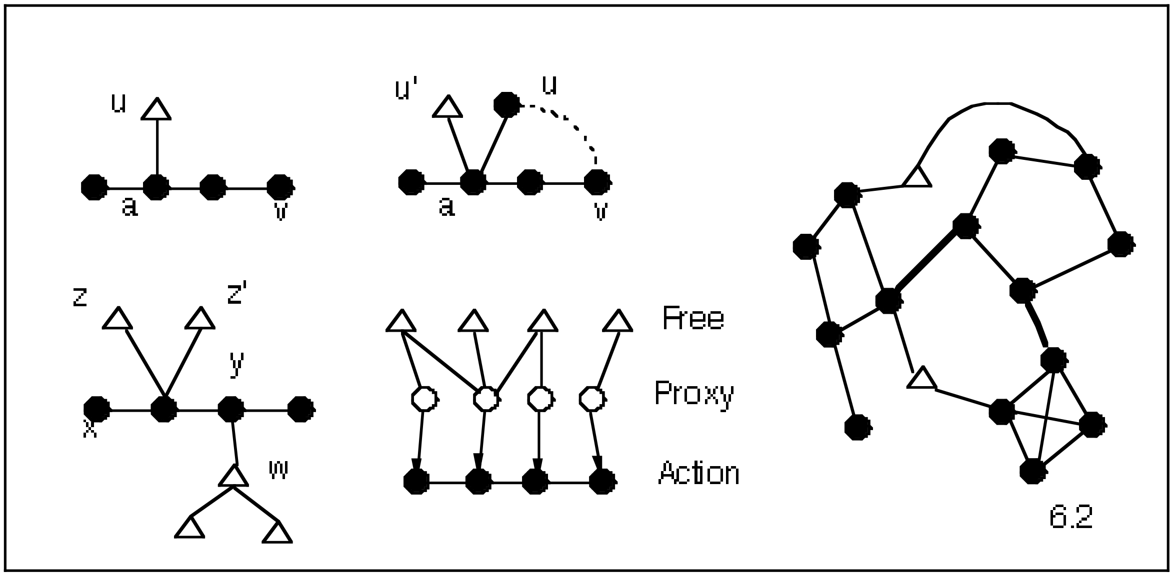

Fig. 6.2. Grafical evolution. The node types are as before, with the lower middle graph indicating the key to node type. The upper left is a 4-node chain of action agents with one free agent addition (u). The adjacent graph is similar, but with (u) now specified as an action node, with a similarity link not only to (a), but also to (v). In addition to single free agents, groups of free agents might be added to an action node (lower left), or to proxies (lower middle.) Free agents might also create loops among a set of action agents (right.) Even with only three types of nodes, challenging network complexities emerge, not only in design, but also in evolution.

6.3 Graphical Evolution

New goals and action states should not totally disrupt mental models that have past successes. Rather, new goals should augment or correct flaws in present models. This implies that:

(i) New goal states (together with their proxies) should be added sparsely, with few initial edges.
(ii) Revised relationships between present action agents should also lead to minimal changes in outcomes (as well as not introducing top cycles).

6.3.1 Node Additions

Referring to the upper left panel of Figure 6.2, let the anigraf form, or model $\boldsymbol{M}_n$, be a simple chain $\boldsymbol{C}_n$-- i.e. the agents or mental organisms have a set of relationships that lie along a path. Then adding a free agent (u) to one of the active agents will introduce a minimum of loops, as compared with an "en masse" augmentation shown at the far right, where now many additional rings appear as subgraphs. A next minimal increment would be to add another free agent (u'), perhaps to the same node (a) or elsewhere. However, adding trees to active agents is also plausible (e.g. "w"), as illustrated at the lower left. Our simple linear chain rapidly expands into a much larger tree. Along the way, the anigraf may recognize new similarity relationships between action agents, creating new edges such as (u,v), which may share a new feature dimension. Behavioral cycles in outcomes now become a larger threat. Also complete subgraphs K_n may be formed, such as if (x) is linked to (y) to create K_3, or the K_4 subgraph shown in the right-most network. Because such subgraphs add little differentiation among the included agents, K_n subgraphs for $n > 3$ should probably be replaced by a new node, characterized by the preferences held in common.

Such informal rules for augmenting any anigraf can be implemented rather easily. One method would be to have a reticular net of free agents as previously described. However, another, perhaps more attractive method, is to convert proxies to action agents who already have similar goals. (See lower middle panel of Fig 6.2). This alternate has the advantage that in the unmodified network, free agents (with zero voting power) can be connected arbitrarily to proxies without changing outcomes. Furthermore, since proxies typically are linked (directionally) to only one action agent, we satisfy the desired three minimality conditions mentioned at the outset. Once a proxy is reassigned as an action agent, its connectivity to its previous host is changed to bidirectional, establishing the desired new similarity relationships, and any previous proxy-proxy relationships would wither.

Clearly, any of the preceding designs, although seemingly simple, pose deep problems for understanding network evolution. Evolving a network of just one type of node may be tractable. But in the case of anigrafs, we already have at least three different types: action agents, proxies, and free agents (with one more type to follow!) These types interact. Even with minimal constraints on these interactions, the theoretical problems clearly become formidable as more and more types of nodes are brought into play.

6.3.2 Edge Additions

Anigraf agents and mental organisms have physical embodiments. These will impose constraints on the network. Specifically, the distance between nodes becomes a significant factor in the ease of establishing connectivity. If a new belief or goal state is added, what insures that this state can indeed be linked to present agents or mental organisms that share similar preferences? If the similarity in goal states is derived from a similarity in the interface device (sensory or motor), then one might expect some physical proximity between the now newly specified agent and its most similar existing action agent. Altready the “proxy” conversion satisfies this condition. Other linkages to free agents may not. The “physical proximity” constraint implies that any sensory interface or effector connections to free agents should parallel those to existing action agents. To facilitate the developmental process, having free agents physically proximate to action agents (and proxies) offers local options for creating subcategories of preferences of present goal sets. One consequence is that, just as the sensory-motor apparatus constrain inferred model dimensions, so will the physical embodiment of the agents constrain realizable paths for model development.

6.3.3 Giant Components

As we continue to refine the beliefs and goals of agents and mental organisms, any small clique of action agents rapidly grows. This growth raises the possibility that many more relationships will be found between members of different cliques. Certain loud sounds may be associated with very bright (lightening) flashes, or temperatures with certain visual colors. Powerful running requires more than just leg motion, but also hip movement and foot exertion for added lift. These in turn affect balance. Similarly, to find food, we must use vision, taste, grasping, locomotion, etc -- a complex, context-sensitive network of quite different agents. So the anigrafs can easily explode well beyond very simple forms like trees, into a tangle of forests, and finally into some large complex network. Three very serious problems emerge:

(i) One tally machine for the millions of agents potentially involved in decision-making seems very unlikely: the cycle time involved (e.g. for neurons) quickly becomes prohibitive.

(ii) Large networks created by combining or linking cliques would be expected to have increasingly large number of edge steps between many of the nodes (and hence agents.) If agents are to share preferences, (a) the distance of their communications would then increase well beyond nearest neighbors. Furthermore, as network size increases, one might expect agents to become more highly connected to other agents, in which case, (b) the set of preference options for any one agent must increase as well.

(iii) Large “random” graphs with “noisy” agents are more likely to lead to cyclic outcomes with no clear winner. (See Appendix 2.)

We can overcome the first part of the second objection of network size (i.e. ii(a)) by placing constraints on the form of very large graphs, and hence the way cliques and free agents are added or merged. If physical constraints permit, one possibility is if graphs evolve into random graphs, where edge connections have a fixed probability. Then for very large graphs, the node distance between any two agents will be at most two edge steps. This “solution”, however, leads us directly into the third objection (iii), as well as the second objection (ii(b)). Agents become overloaded on the number of preference options they must be capable of sharing. And finally, this solution does not solve the problem of tally time.

A second, more attractive option is to have graphs evolve into a multi-scale structure, having a hierarchical “tree-like”form, with cliques of agents having a wide range of sizes (a fractal distribution), and with sparse links between cliques (see Kasturirangan, 2002.) Mental organisms at each scale would then reach their own consensus, hopefully with the tallies be conducted simultaneously, or at least in harmony. (See recent work by Bruice & Cowan, 2007, relevant to tally cycles conducted at different spatial scales.) But have our babble of daemons now returned, with each clique shouting for control? Can we control which mental organisms should dominate at any one time?

Spiral Forms (courtesy of P. Gunkel,)

Anigraf7: Planners
event sequencing

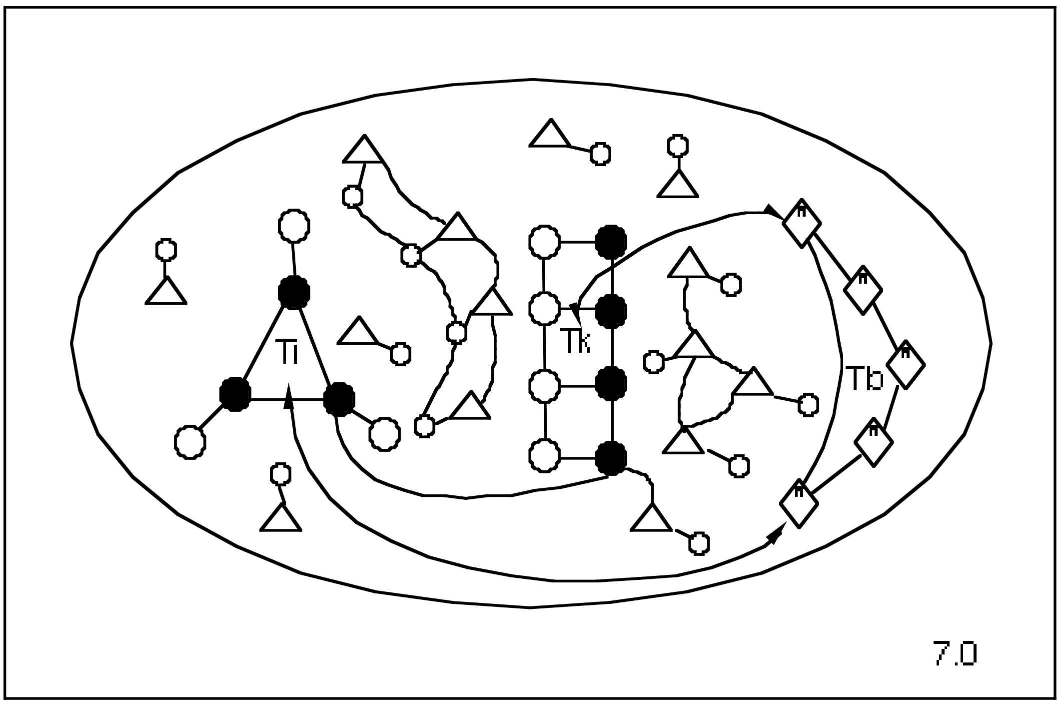

Fig. 7.0. Depiction of a network with three active mental organisms. Two have action agents (solid circles), proxies (open circles) and individual tally machines, T_i, T_k. The third is a broker consisting of five event states (diamonds), and tally machine T_b. Note that some of the event states are linked to one of the two different cliques of action agents that more directly initiate actions such as limb movements. In the background we see several sets of free agents (open triangles) connected to potential proxies.

10 Aug 07

7.0 Event Sequencing

If anigrafs organize multi-scale networks into cliques of sub-graphs, not only do we have pandemonium with many sets of mental organisms clamoring for attention, but also we have the possibility of multiple, conflicting demands

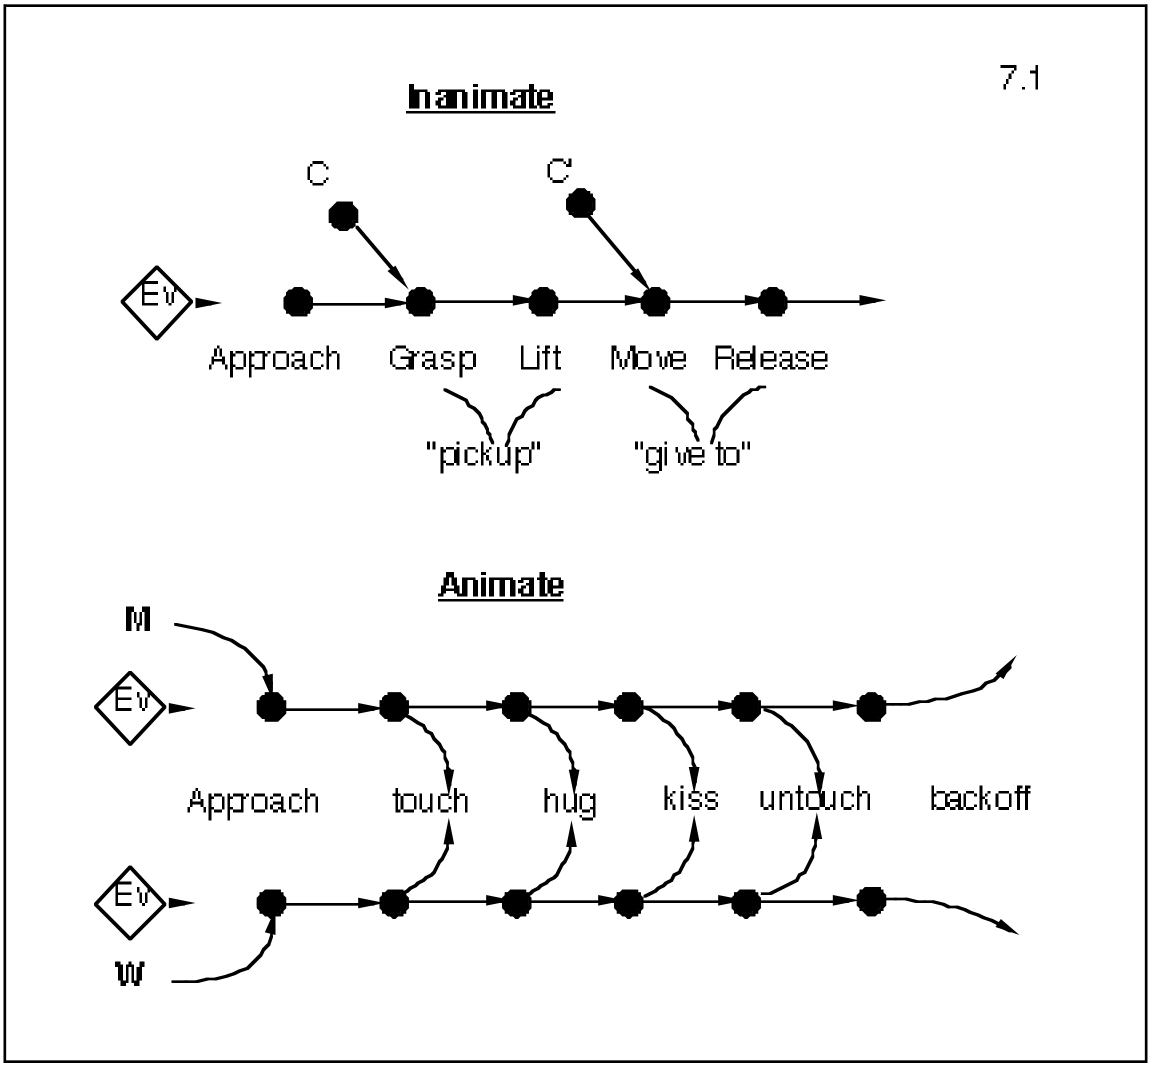

Fig. 7.1: Two sequences of events (E_v). In the inanimate example, the anigraf's task is to sequence a set of actions to pickup and release an inanimate object. In the lower "ladder-like" graph, two anigrafs need to program a sequence of coordinated events.

on the same motor system. If each sub-graph clique is conducting its own tally at its own rate, then any coordination among the active mental organisms becomes unlikely. For example, let the task be "to get food". Then, as shown by the upper graph in Fig.7.1, there is an obvious sequence: look, approach, grasp, eat, etc. If the task is to move a boulder, it must first be loosened, then rolled. Similarly, an activity that requires coordination between two anigrafs, will require a proper sequence of steps before a goal can be consummated. We need agents and mental organisms that set up these sequences.

The most obvious solution is to have a collection of high-level mental organisms, each of whom control a set of specific actions appropriate for every desired sequence. This scheme is rejected on two grounds: (1) even for simple creatures, the number of different sequences that need to be constructed becomes prohibitive, and (2) the organisms controlling the sequencing would have to draw upon action agents who are members of various cliques, very likely at quite different physical locations, resulting dense networks with very complex wiring diagrams. Furthermore, the local tallies within each clique would then be replaced by a different set of tallies conducted by the particular mental organism that encoded each sequence. In effect, there would be one large jungle of smaller graphs with a host of different tallies for all, each hopefully able to carry out the requisite n-cycles needed for each sequence.

To avoid these objections, we need a process that leaves some control of the sequence to individual cliques of action agents, allowing other mental organisms to focus on the choreography of event sequencing. Let us identify these higher-level mental organisms as “brokers”. Specifically, if the brokers controlled only the category of actions and the order in which cliques of action agents are chosen, then the connectivity is to appropriate cliques and not to all individual action agents within each clique. The brokers would thus bring different categories of agents together in the desired sequence, but allow the cliques to pick that action most appropriate for the context. If the sequencing is correct, then inputs from the current world state (and the creature’s own body) will properly drive the tallies when each clique is called upon to vote. The world thus acts like an external memory with physical constraints that partially control “what to do next.”

7.1 Depiction

Fig 7.2 shows a possible setup. The aim is to highlight some of the conditions needed, but not to present a formal model. Let us build on three mental organisms, $\{A_i\}$, $\{B_i\}$, $\{C_i\}$, each with different tally machines (T_a, T_b, T_c) and different graphical models linking their different sets of action agents (solid circles). Each clique initiates a different kind of action. Clique A, for

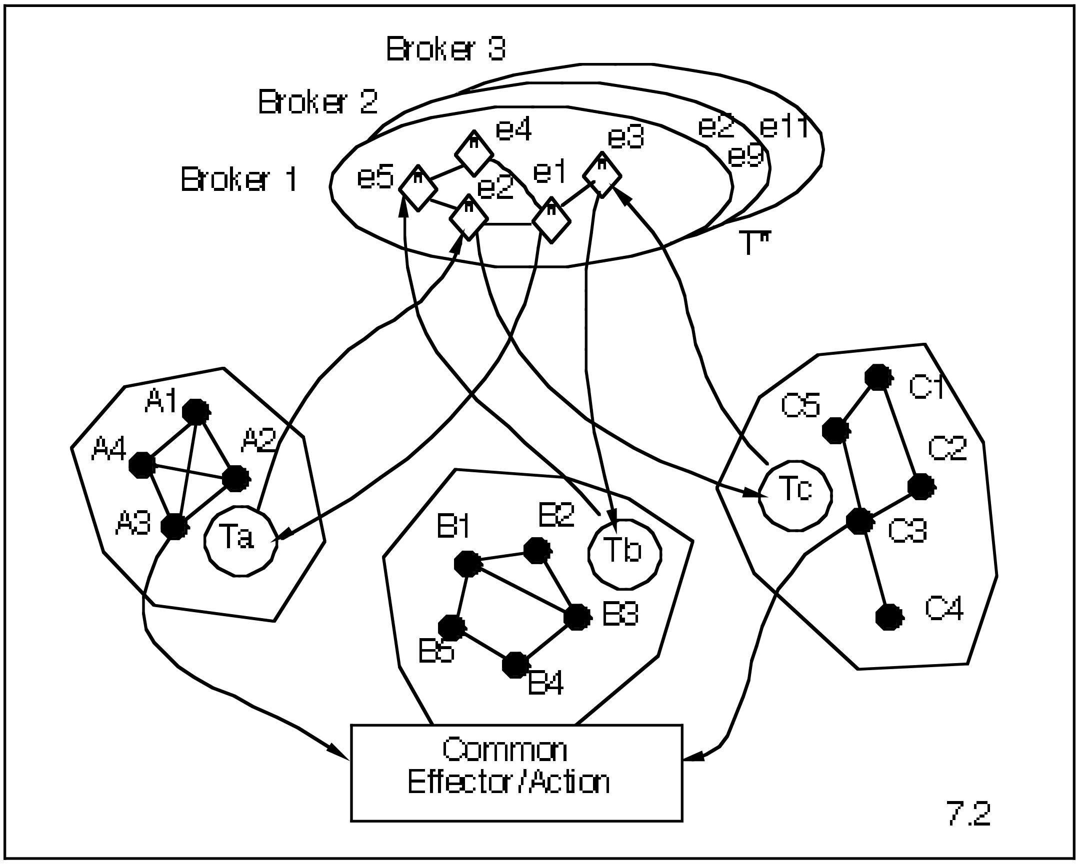

Fig. 7.2: At the top are three different Brokers, each with a set of events whose sequence they control. (Note that one event may be common to two or more brokers). The tally T* dictates the choice of broker. Each broker chooses and regulates a sequence of actions, according to the outcome of its own tally. Below the broker ellipses are three different cliques of action agents (A, B, C). Each clique is reciprocally linked to one or more brokers, but only through their tally device, $T_{i,}$, not via individual action agents (solid circles.) Back projections from the T_i's are used to trigger the next event in the selected cycle.

example, may govern gait type (walk forward, retreat, etc.); clique B possibly how things are grasped (pickup, release, push); and clique C perhaps whether something should be eaten or simply bit into or poked. Note that for a given set of categories, different sequences are possible. If the goal is to eat an object, the clique sequence would be A -> B -> C, whereas if food is offered by another anigraf, the sequence might be B -> C -> A. The role of the brokers is to decide which sequence of cliques is appropriate. To accomplish this task each broker has a collection of possible event types e_i, such as grasp, look, move, etc. which are related by the broker's model. At the top of Fig 7.2 we show three ellipses that indicate such mental models of three brokers. The event types e_i are indicated by starred diamonds. Note that the same event type may appear in different broker models. There is also a global tally T* conducted among all the brokers to choose that broker who will choose and control the sequence of cliques of action agents. Not shown are the graphical relations between the goals of the various brokers. These relations may be dictated in part when brokers share the same event e_j, or when their event graph structures are similar.

7.2 Innate Releasing Mechanisms

It is important to realize that the hierarchical control of actions by the brokers does not guarantee that the desired set of actions take place. Brokers only control the sequence of the categories of actions, not specific actions. (If brokers were to control specific actions, the event sequence would have the form A_1 -> C_3 -> B_5, for example.) At best, brokers can only specify the categorical sequence, namely A -> C -> B, leaving the choice of specific actions up to the members of each clique. As mentioned, the environment in which the anigraf is currently embedded plays the key role in choosing the correct action. The situation is much like when a cook prepares a dish: the ingredients are laid out on the counter, and once a sequence of steps is initiated, certain operations become obvious, whereas others are clearly inadmissible.

In lower animals, many behaviors are ritualistic, where one act sets up the preconditions for the next. The building of a spider's web, or the burying of a nut, or simply moving a boulder all entail steps where one act changes the external state that in turn allows the next act to proceed. This is especially clear

Brains and Modularity

Anigrafs are abstract entities, like numbers or ideas. Their physical correlates are loosely specified principally for pedagogic clarity. An exception is the tally machine, which is the physical device that provides a link between mind and brain. Otherwise, nodes (or vertices) could be human beings in a community; nation states; perhaps micro-scale entities; and, of course, neurons or various neural modules which may dominate the reader's interpretations. Similarly, the communication links in anigrafs (i.e. the edges), need not be axons or other wire-like channels, but can include a host of possible wireless links, every bonds, field potentials, or perhaps even some quantum-like interaction. All such assignment of physical forms to anigrafs will affect their designs. For example, a tally requires a physical device; and each tally is associated with and identifies a particular anigraf. When different anigrafs need to be coordinated, then so do their tallies. Communication rates impose strong constraints on synchronization and sequencing. The time-constants of sensory-motor devices also become a factor, as does network size. Modularity has been favored as one approach to choreographing a symphony of behaviors, rather than attempting to build just one extremely large anigraf with one tally for all mental organisms and action agents. Clearly there is an interplay between structure and conceptual design. In the spatial dimension we have another constraint related to modular constructions. If edges are equated to wired links between nodes, rather than choosing wireless links now common in social communications, then circuit designs limit anigraf forms. Suppose that we require that no edges overlap, touching one another, and that all circuit components must be laid out on 2D sheets. Then many anigraf forms are excluded. Alternately, one might impose limits on the number of 2D surfaces onto which a graph could be embedded (i.e. the "thickness" of a graph.) Or, perhaps, structural factors could limit the genus of graphs (the number of "handles" required on a sphere to prevent edge crossings.) Clearly, if structural considerations come into play, then modular designs can be very beneficial. But at the same time, structure will be imposing constraints on conceptual models. An anigraf's view of Nature will again become limited.

in mating behaviors, where consummation requires first various displays that arouse interest, then contact, and finally intercourse. When such coordinated sequences occur between two creatures, intentions need to be read correctly (unless fully automated), and hence brokers, like all other agents, can be expected to have their proxy counterparts, as well as free agents who will provide the opportunity to learn new sequences in new contexts.

7.3 Event Sequence Control

Each broker is presumed to have control over a set of events that carry out a specific class of actions appropriate for some goal. The broker that gains control is chosen in the usual manner by a tally T* among brokers, given the weights placed upon the different goals at that time. Once a broker is chosen, however, we have left unspecified just how the particular event sequence is determined.

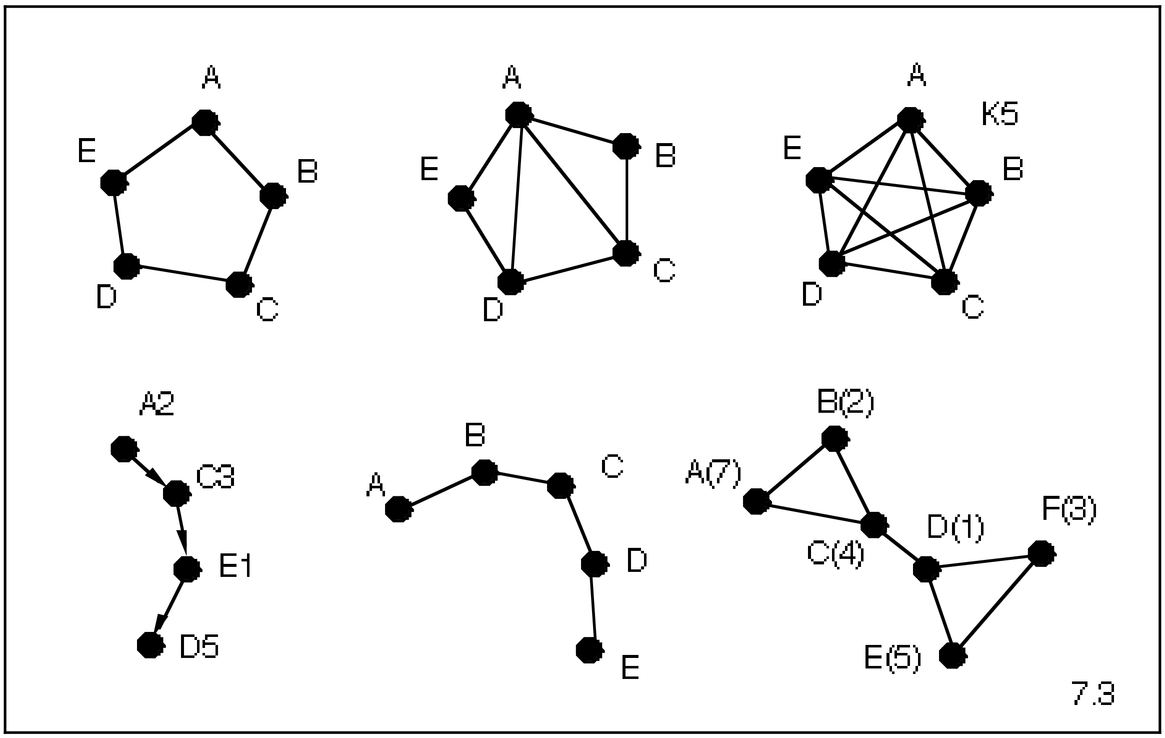

Fig. 7.3: A set of simple mental models for controlling event sequences. At the lower left, the sequence follows the directed edges. For ring graphs, the sequence can easily be altered by a phase shift in weights. The remaining graphs can also be used for controlling brokered sequences, but may require modifications of the tally procedure. (See text.)

7.3.1 Choreographed Cycles

The most obvious method for event sequence control is analogous to an innate releasing mechanism where for each broker has only one pre-programmed sequence that runs when initiated. For example, as illustrated at the lower left of Fig 7.3, the directed sequence A2 -> C3 -> E1 -> D5 might ensue regardless of weights on the action cliques. However, to give our brokers more flexibility, we could have the sequence determined by a social cycle among the possible events, such as we did previously in a dance. In this case, different sets of weights could generate different event sequences. For example, if a broker had arranged events in a ring with an odd number of nodes (events), then various triples of event sequences become possible given the proper set of weights. For small rings, choosing between different cycles becomes trivial. For the bi-directional pentagon in the upper left of Fig 7.3, the only difference between cycle ACD and BDE is that the weight set has been shifted clockwise by one vertex.

7.3.2 Strictly by Weight

Let's suppose that the initial step in a sequence finally gets the anigraf's approval, and has the maximum weight, and the second step has the next greatest weight, etc. Control of a sequence maps directly to the ordering of weights. In this case, the broker's anigraf is very simple: the relation between events is the complete graph K_n (K_5 is illustrated.) Because all event nodes are connected to each other, in any pairwise comparison between two nodes, all remaining nodes are neighbors to the combatants. Hence they will be indifferent in their preference. The node with maximum weight will then win. Once the weight on this node is set to zero, the node with the next largest weight will win, etc. (Note that if the final event is given the most weight, then one would use a reverse Condorcet tally.)

A variant of this procedure would be when only one node covers all others, as shown in the top middle of Fig. 7.3. Event category A now wins the first vote roughly 70% of the time, because it receives support from all other nodes. But now what if the weight on A is zero? A will still win 16% of the time. Hence an ordered sequence of weights is not guaranteed to produce a comparable sequence of node choices.

7.3.3 Constrained Sequences

The complete and the covered graph are among the simplest examples illustrating the interplay between graphical form and node elimination. Consider the bow-tie shown at the lower right of Fig 7.3. Let the weights on nodes (A, B, ….F) be respectively (7, 2, 4, 1, 5, 3.) When a tally is taken between any element of the triangle ABC vs DEF, the winner will be that triangle having the greatest weight sum (or lowest if reversed Condorcet.) Hence the first winner is A, who beats all its neighbors, as well as D, E, F who together can muster only 1+3+5=9 votes against A's total of 13. But once clique A is triggered, and its weight drops to zero, then the voting power will pass to the right triangle DEF. E now becomes the new winner. Continuing this process for the weights shown will elicit the sequence A -> E -> C -> F -> B -> D. This particular graph will favor alterations between the two triangles, that might correspond, say, to a sequence of various types of symmetric body motions each followed by different limb motions.

7.4 Maxed Out

Hierarchical anigrafs with brokers controlling event sequences has been stretched to a level of complexity that raises many questions, and impedes further lucid presentations. Additional proposals for embodied anigraf designs can only make the construction more opaque, endangering a simple picture of creatures whose minds are social networks. Nevertheless, there is one more insight that can be gleaned from our brokers. What if the set is not embodied in one physical entity, but rather are seen as a collection of different peoples, each with their own game plan? Together, they would form a society in the popular sense. What kinds of behaviors might be expected?

Anigraf8: Alliances
coordinating diversity

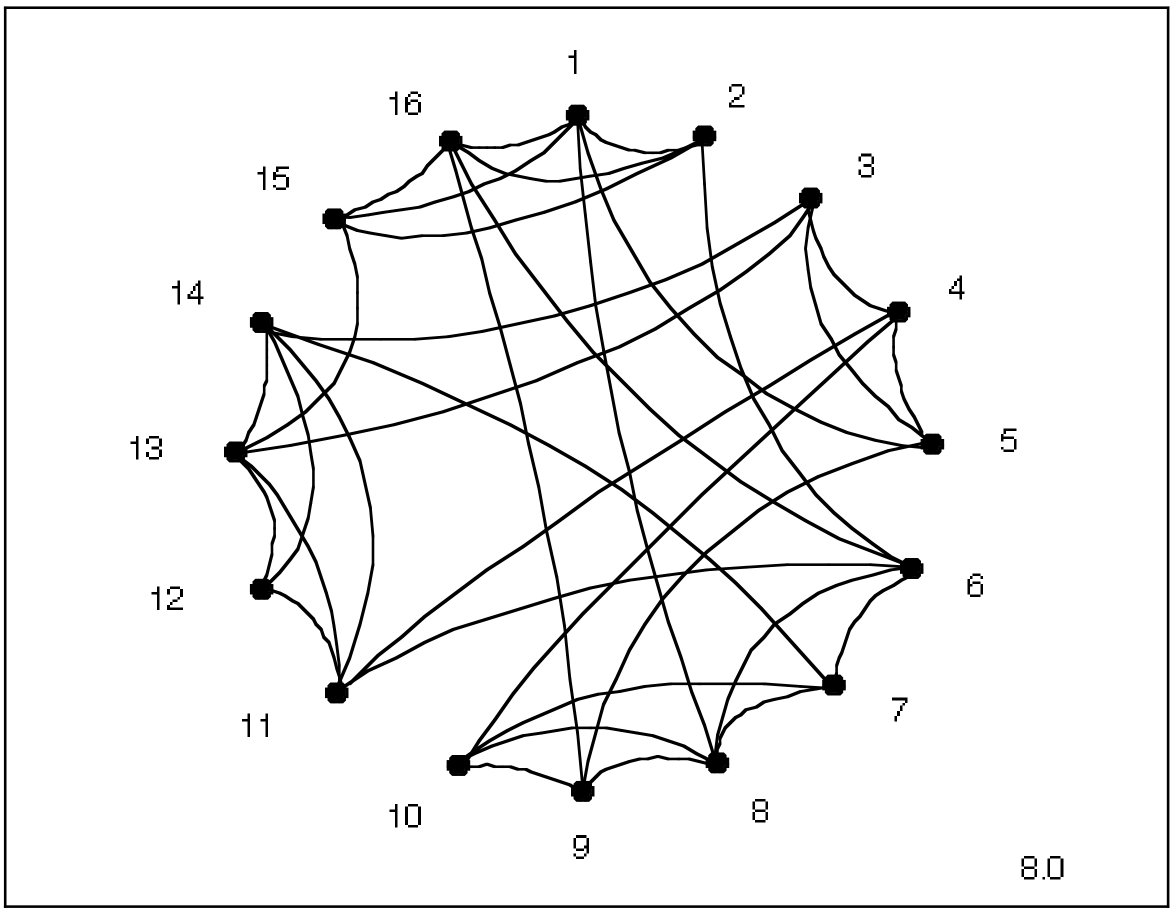

Fig. 8.0: Schematic of a small network of friends. When will one group of friends become aligned with another group?

11 Aug 07

8.0 Heterogeneity

Consider a society of identical anigrafs. With such homogeneity, we can expect accurate, intelligible communications among members of the society, but at the sacrifice of the ability of the group to see the world from different vantage points. Behaviors will also be very predictable. Clearly, the flexibility and adaptability of a society depends upon its members possessing a range of talents to execute a host of different tasks. Successful models for one situation may be quite inappropriate for another. Fortunately, given even minimal environmental pressures, we find that offspring are not carbon copies of their parents. As models grow to deal with new events and challenges in the world, the offspring's view continues to depart more and more from the parental path. Diversity evolves, even in the presence of a common core. We have, then, the evolution of a colony with shared goals, but with members of different talents. Examples include insect societies, football teams, a jazz combo. Although their expertise may differ, the members are united as one entity. How different may any two members be, and still experience this unity? The answer sets the stage for studying how a population of anigrafs with different models may be able to form productive alliances.

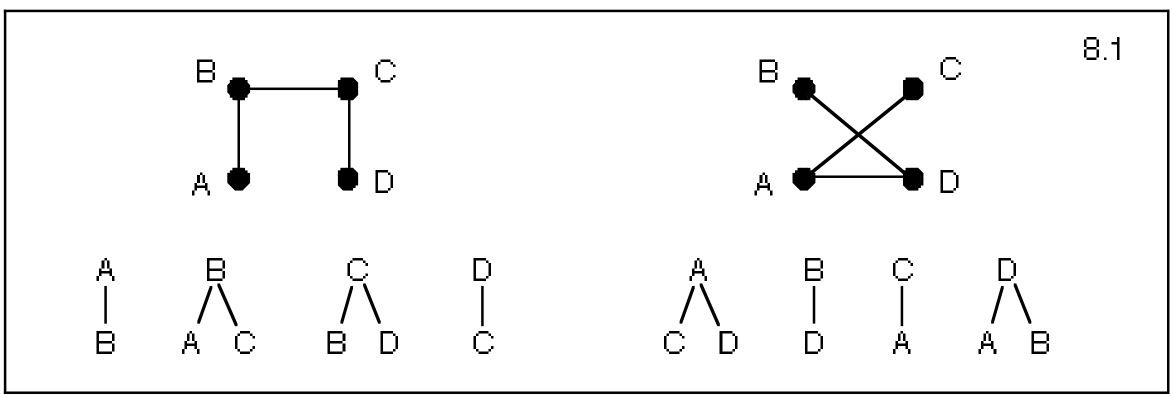

Fig. 8.1: two complementary graphs and their preference orders.

8.1 Similarity and Dissimilarity.

In anigraf world, similarity and dissimilarity are tied to graphical forms, or more precisely, the derived preference orderings. These preference orderings reflect goals and choices related by the mental organisms that comprise anigrafs. If two anigrafs have mental organisms with the same set of goals, then these anigrafs will be identical only if their graphical forms are the same. An opposite extreme would be two anigrafs with the same goals and choices, but with preference orderings derived from a very different model. To illustrate, Fig 8.1 shows two graphical models that are complementary: each are subgraphs of K_4 and together form the complete graph. None of the preference orderings agree. Hence, even given the same set of weights on alternatives, the Condorcet choices would typically differ. Our aim is to have a diversity in anigraf forms, yet still have enough similarity among members of the group so that these differences do not prohibit coordinated actions. This is the essence of building an integrated team of players.

To capture our intuition about similarity and diversity (or equivalently, dissimilarity), we can measure the fraction of agents or mental organisms in both anigrafs who have identical preference orders. This measure automatically takes into account dissimilarities introduced by differences in goals.

Definition: *The **similarity S(kd)** between two anigrafs having n and m vertices is the fraction of all (n+m) vertices that have the same preference orderings, where "**kd** " is the depth of the preference orderings.*

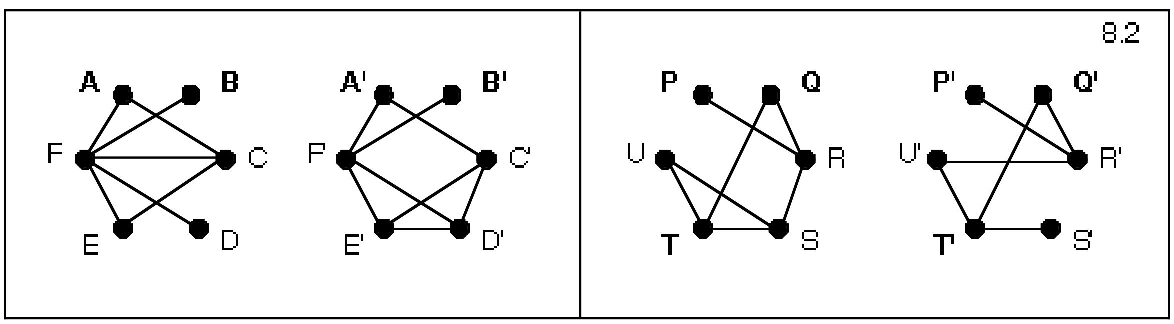

Fig. 8.2: Examples of the anigraf similarity measure, with kd=2. On the left, the similarity is 1/3; the right pair have similarity 1/2.

The patterns in Fig 8.2 exhibit two pairs of anigrafs. The pairs A, A', P, P', etc., have identical goals, and bold type indicates which pairs of vertices have the same preference orderings. To generate these anigrafs, the edge set of the left member was chosen with probability 1/2. A core for the second anigraf was then created by freezing vertices A, B and P, Q, keeping their same edge sets. The remaining four vertices were then assigned edges with probability 1/2 using the reduced edge set. For the anigrafs with vertices {A, B.....F}, none of the remaining four vertices {C, D, E, F} had the same edges. Hence for kd=2 (neighbors only), the preference orders for A and B are identical to A', B' and the similarity *S(2)* = 1/3. For the {P, Q....U} pairs, vertices P, P' and Q, Q' have the same edges, and so do vertices T,T'. Hence *S(2)* = 1/2. Structural similarity for anigrafs is thus dictated by vertex similarities and only indirectly by the edge sets.

8.2 Similarity vs Outcomes

As the similarity between two anigraf decreases, so will the percent of outcomes that are the same. To explore this relationship, generate a set of random graphs $\{G_i\}$ with n vertices and with edges chosen with probability one-half. For each of these graphs we now replace a fraction, f, of the edges to create a new set of graphs $\{H_i\}$. For large n, we find that the chance of G_i and H_i having the same Condorcet winner is proportional to the similarity between H_i and G_i. This relation is shown by the diagonal line in Fig. 8.3.

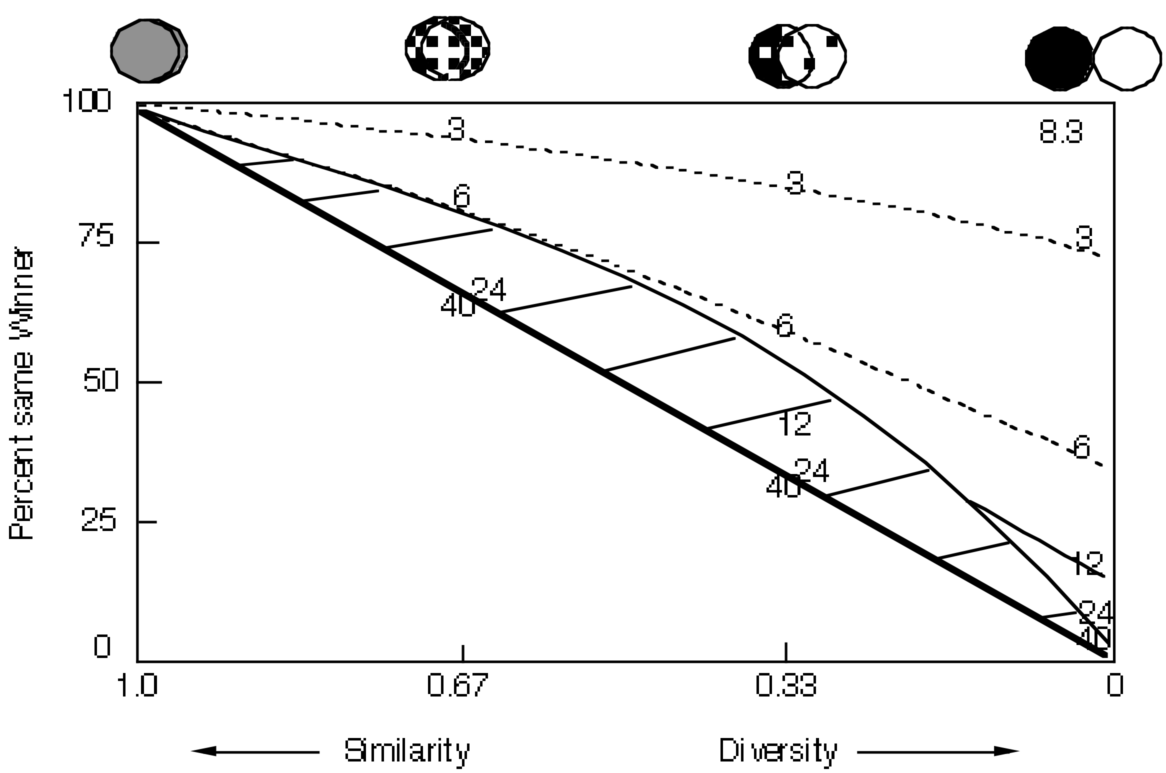

Fig. 8.3: The similarity between two anigrafs, each with edges chosen with probability one-half. The preference orderings are limited to neighbors (kd=2.) The labels indicate the size of the graphs.

One might ask whether all tally procedures will generate a result similar to Fig 8.3. The answer is no. If the vertex with maximum weight is always chosen as the winner, then all curves will be flat at 100% regardless of their similarity. The Condoret tally is special in that it incorporates information about second or higher order preferences, which was used as the basis for defining similarity. Knowledge depth (kd) now becomes an important parameter in relating similarities and outcomes. If the knowledge depth is increased, the results of Fig 8.3 will change, especially for small graphs. The reason is because deeper preference rankings will include relations for vertices previously chosen to have fixed edges to their neighbors. Consider the right pair of anigrafs in Fig 8.2. If knowledge depth is increased from one to full for the left anigraf, then the preference order for P changes from [P>R>Q,S,T,U] to [P>R>Q~S>T~U]. Hence the Condorcet (and also Borda*) tallies will be change, subject to the particular way the second anigraf is augmented. The exception will be for graphs with diameter two (which include the set of asymptotically large random graphs.)

8.3 Coordination and Diversity

Our objective is to create a diverse group of anigrafs that still have some unity through shared goals and preferences. But figure 8.3 presents a dilemma: as graphical forms become less similar and more diverse, the odds for a pair of random anigrafs having the same winner decrease proportionately. How can we optimize diversity and similarity at the same time ?

Although the linear relation between similarity and common outcomes implies only one underlying variable, this is not the case. There is another variable obscured by the deceivingly simple figure. Absent is just how tightly the percent of common winners are distributed about the mean for any given similarity. For example, a set of quintuplet anigrafs, or a set of K_n anigrafs, will always agree on winners, and hence their variance is zero. Likewise, there will be zero variance for a collection of (large) random anigrafs, where the probability of any pair having the same winner is zero. Between these two extremes, there will be a distribution of the percent of common winners for any similarity index. This relation is shown by the cross-hatched envelope in Fig 8.3 for anigrafs of size 24. We now use this variation to advantage.

Let us repeat the initial experiment where G_i and H_i have different graphical forms, but now also compare winners with a third or fourth variation, H_i J_i or K_i of G_i, where all variations have the same similarity measure with respect to G_i. Using Fig 8.3, we can now estimate the probability that at least one of the pairings with G_i, namely H_i or J_i, etc. will have the same winner. This pairing probability gives an indication of the strength of group unity. Multiplying this probability-of-unity times the dissimilarity of form, namely (1 - *S(2)*), gives a measure of "coordinated-diversity." Maximizing this product yields the optimal balance between coordination, or group unity, and diversity.

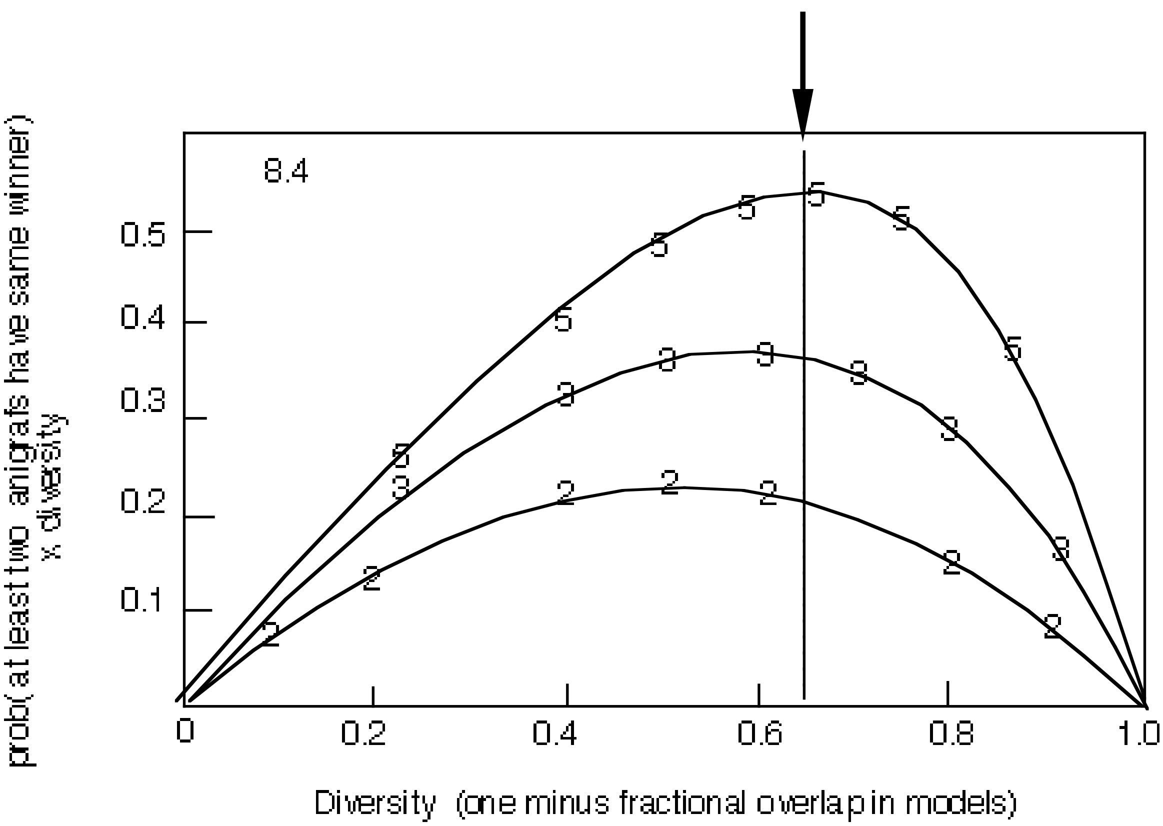

Fig 8.4: Plot of coordinated diversity for groups of 2, 3, and 5 anigrafs, Each group has 12 nodes, with the edges chosen with probability one-half from a uniform distribution.

Fig. 8.4 shows this "coordinated-diversity" product for small groups of size 2, 3, and 5. The maximum moves towards a two-thirds dissimilarity of form. The simulations used 12-node random graphs with edge probability of one-half (dbl-ck.) Thus an estimate for the optimal diversity among members of an anigraf population is when the similarity between pairs is about one-third.

8.4 Clique Formation

Within any population of anigrafs, benefits to some can be gained by forming subgroups where members agree to vote as a block. Obviously, members of such groups will resemble one another and will have strong similarities in preferences and choice of winners. What role will the anigraf form play in the likelihood that a set of anigrafs will be able to share common preferences?

Definition: *A* ***clique*** *of size m is a set of anigrafs of size n where each of the m members shares at least one preference relation with at least one other member of the clique*.

In other words, if any two members of a clique are chosen, there will be at least one goal for each anigraf that is the same for both, and, for knowledge depth = 2, will have the same alternate goals or choices (i.e. the neighbors to that vertex in the anigraf are the same.)

Consider now random anigrafs of size 12 with edge probability one-half. From Fig 8.3 we know that there is roughly a 10% chance that any two anigrafs will have on average one agent or mental organism with the same preference orders (to knowledge depth two.) In such a population with a small likelihood of even one shared vertex between members, the clique size will be small (relative to the population size), and conversely, the number of cliques will be large. However, if we were to increase the edge probability toward one, then the anigrafs will become increasingly more similar, eventually all being identical, forming one large clique. This effect of edge probability on expected clique size is shown in Fig 8.5 for a population of 20 different anigrafs.

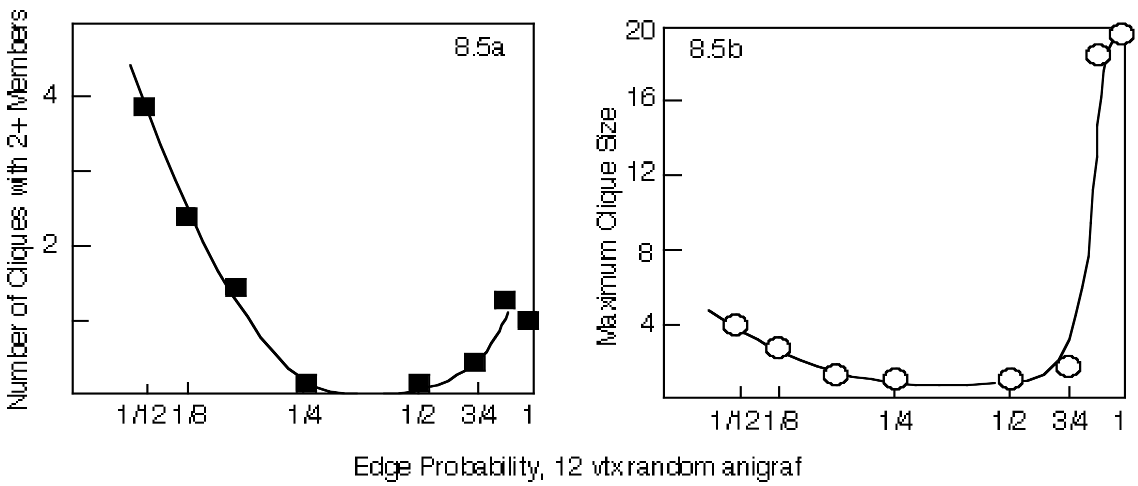

Fig. 8.5: Both graphs describe partitions of a set of 20 *different* anigrafs. Any clique requires that each anigraf have at least one preference ordering that is identical to another member of the clique (kd=2.) Obviously, as shown in the right panel, if the edge probability is very high, most of the 20 angrafs belong to the same clique. Note the surprising trend at the left of this panel, however. In the left panel, when the edge probability approaches 1, there will be only one type of anigraf, and hence a degenerate clique of one. See text for explanations.

In the right panel, for edge probability one-half, almost all individual anigrafs are isolated, and the maximum clique size is very small. As edge probability goes to one, clique size increases to the size of the population of 20. Along with this trend, as shown in the left panel, the average number of cliques with at least two members will increase, except at the limiting value when the network is completely self-connected.

The surprise is when anigraf edges become sparse, with edge probability going toward zero. Now maximum clique size increases again and more and more cliques appear with at least two members. This multi-modal distribution is a key to the formation of Social Networks: If a unified population is to consist of dissimilar "random" anigrafs, then the division of this population into cliques of any significant size will require sparse anigraf designs. But sparse anigrafs begin to resemble trees.

8.5 Social Networks

We are now in a position to explore some aspects of anigraf social networks. Such networks examine the similarities between anigrafs of different sizes and types. Roughly speaking, a social network is a network of cliques, showing how different species can remain compatible through similarities in structure. When anigraf forms resemble tree graphs, the similarity will be via a sub-tree, for example whether a vertex has three leaves or four, or ends in a star-burst, etc. The main result is that populations of sparse anigrafs of different sizes will tend to form scale-free networks. (Watts, 1998; Barabasi, 1999, Kasturirangan, 2002.)

Definition: *A **social network** is a set of loosely connected cliques, with each member of the population belonging only to one clique.*

By loosely connected, we mean that the density of similar vertices between cliques is much less than the density of similar vertices between members within a clique. Again, we fix the knowledge depth of the anigrafs at one, considering only neighbors to vertices when evaluating similarity.

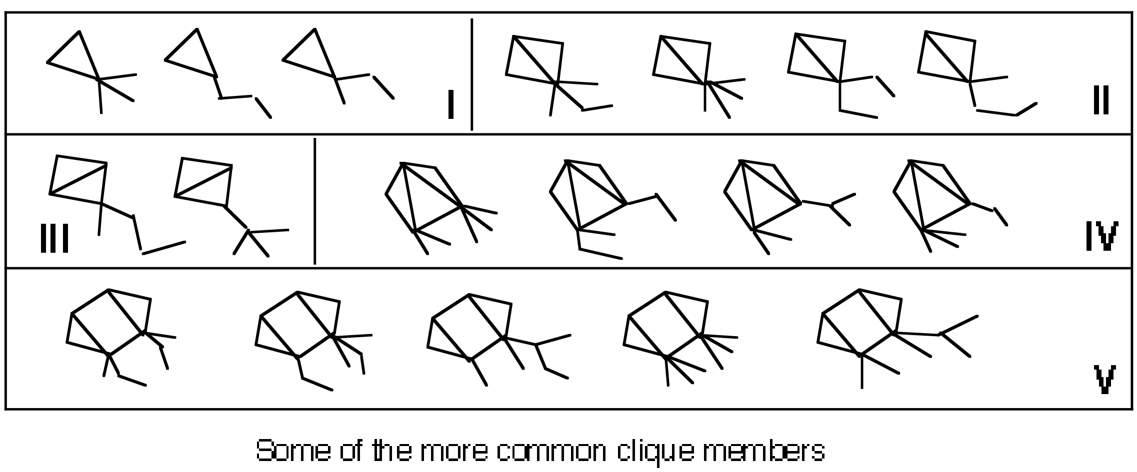

Fig. 8.6: Five different species of anigrafs. Each anigraf model has two parts: a planar polygon composed of one or two triangles and a square (as in type V); plus a tree-like part linked to one or two nodes of the polygonal model.

Let us now generate a population of 50 labeled anigrafs having 6, 8, 10 and 12 nodes. Each of these has one fixed subgraph, such as a triangle, a square with one diagonal, a pentagon with one vertex of degree four, etc. (See Fig. 8.6 for examples.) Attached to this fixed subgraph or "body" are limbs and appendages. These attachments are generated as random trees, always rooted at the same vertices of the "body." Both the fixed "body" subgraphs and their attached random trees have the same number of vertices, namely 3, 4, 5, and 6 respectively for the different size anigrafs. The design of the "body" subgraphs insures that cliques will be formed among among the different species of anigrafs; the random tree attachments provide the basis for similarities between different species, thus forming the basis for a social network.

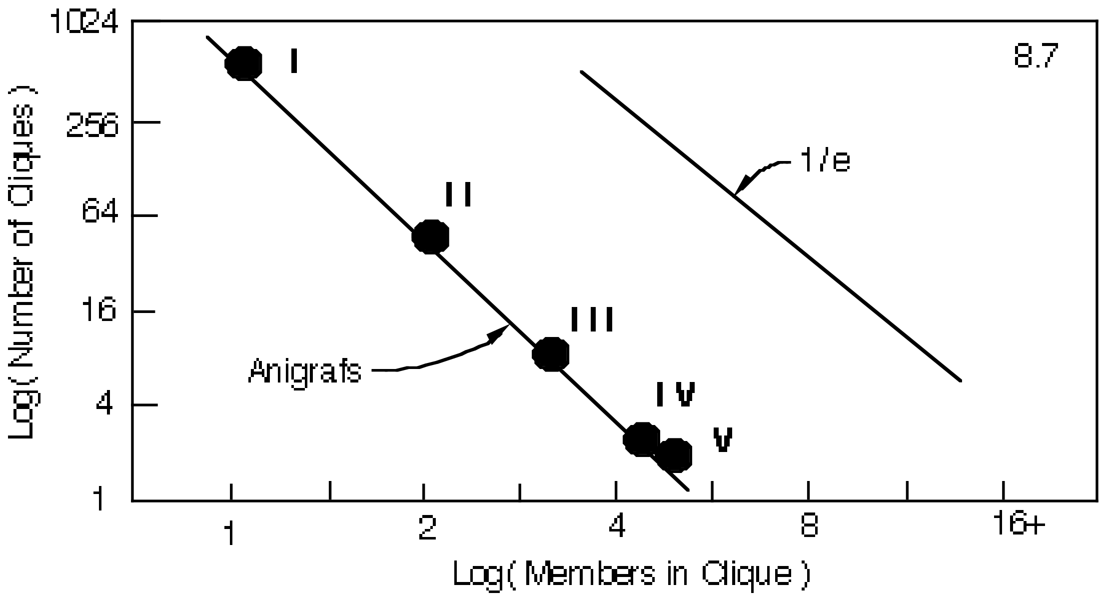

Fig. 8.7: Clique membership size for anigraf species illustrated in the previous figure.

Figure 8.7 shows the number of cliques of any anigraf size that are formed when several populations are generated, and when we require four similar vertices among at least two members of a clique. There is a large number of small cliques, and few very large cliques. It is the tree-like form of the limbs of the anigrafs that lead to this relation. Vertices of degree one are easiest to link, then those of degree two, etc. Hence, because their tree is larger, larger anigrafs will be more likely to have vertices of degree one and consequently will find more clique members. This leads to a 1/e slope of the relation between clique size and number, where e = 2.718, the basis for natural logarithms.

Although each clique by itself has a social structure reflected in its clustering properties (on average 80% of that of the comparable complete graph), this is not the social network of interest here. Different species can not exist independently of others. How will a society of anigrafs of different types and sizes relate to one another? For our primitive anigraf collection, one answer lies in how the various cliques are linked and interconnected. If at least two similar vertices are required to link anigrafs of any size and type, then the most common affinities will be among cliques with the smallest anigrafs. This follows because there are very few different tree-designs for their limbs. In contrast, larger anigrafs with six or more vertex trees will have many forms for their limbs. Most of these forms will be quite different from those of the very small anigrafs with only three-vertex trees. Hence linkages in the larger social network will favor anigrafs of the same size, or anigrafs of only slightly different sizes. Appendix 8 shows these relations in more detail. The consequence is that the density of similarity links within a clique will narrow the ability of clique members to align with others of either larger or smaller size – or to anigrafs with different basic designs. At a very large social scale, this issue has implications for the stability of ideological forms adopted by nation states and the nature of competition among global orders (Hayek, 1952.) On the very small scale, the problem lies at the heart of how easily one anigraf can relate to and understand the models of another, thereby influencing behaviors.

Anigraf9: Mind Games

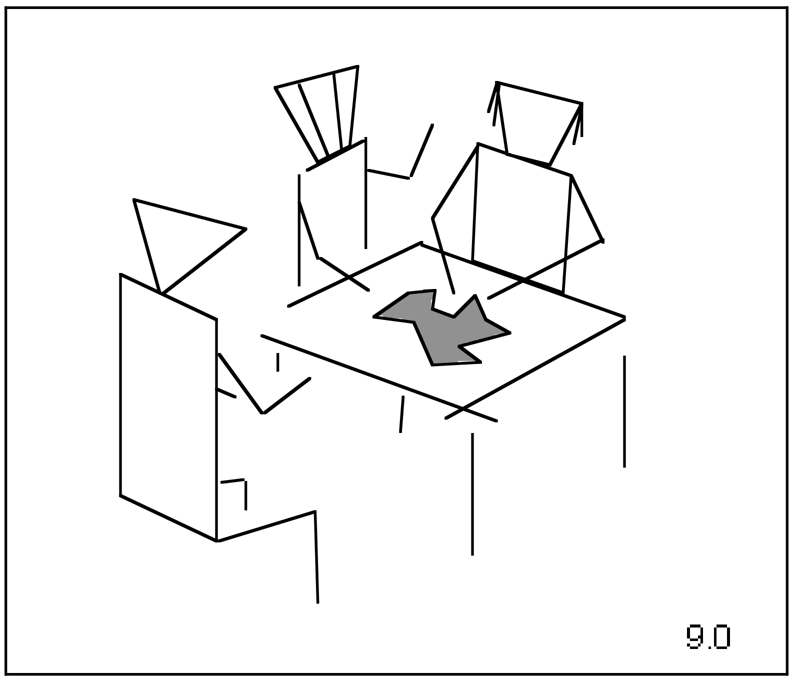

Fig. 9.0: The poker players

10 Aug 07

9.0 Theory of Mind

Powerful alliances between anigrafs consolidate the diverse abilities of its members into one cohesive unit. Such consolidations increase the dominance of the species. In the cognitive world of anigrafs, the greater the diversity of models in an alliance, the greater the likelihood of predicting intentions of competitive entities. But, as yet, there are no anigrafs that explicitly examine models of others to determine a strategy for optimizing social encounters. Specifically, we need anigrafs that have the ability to “read minds.” If one anigraf could deduce another’s internal model and hence that anigraf’s intentions, then there is not only the possibility for greater clarity in communication, but also the chance for strategic manipulation – i.e. mind games. A new level of intelligence and social awareness emerges.

9.1 The Analyst

Let us focus on one anigraf "P" who is endowed with special analytical powers. Think of him as a psychiatrist. His task is to observe the behavior of other anigrafs, and to infer their mental models. P’s current target is Q. We begin with the assumption that Q’s choices are honest, with no attempt to mislead others. Also, let P and Q have the same numbers and types of constituent mental organisms, which carry out the same set of possible actions. So P and Q differ only in how these actions are seen as related.

After many observations of Q’s behavior when different weights are placed on alternatives, P will obtain a very good estimate of the apriori probability distribution for Q's winners. If these priors are flat, then P knows that Q’s model must be a regular graph, such as a ring, or a ring with chords such that all vertices have the same degree. If the alternative with the maximum weight always wins, then P knows that Q believes everything is similar to everything else, having as an internal model the complete graph K_n. If one vertex almost always wins, then that alternative very likely covers the other alternatives in the anigraf model. The priors on the actions Q takes thus provides important clues to Q’s internal model.

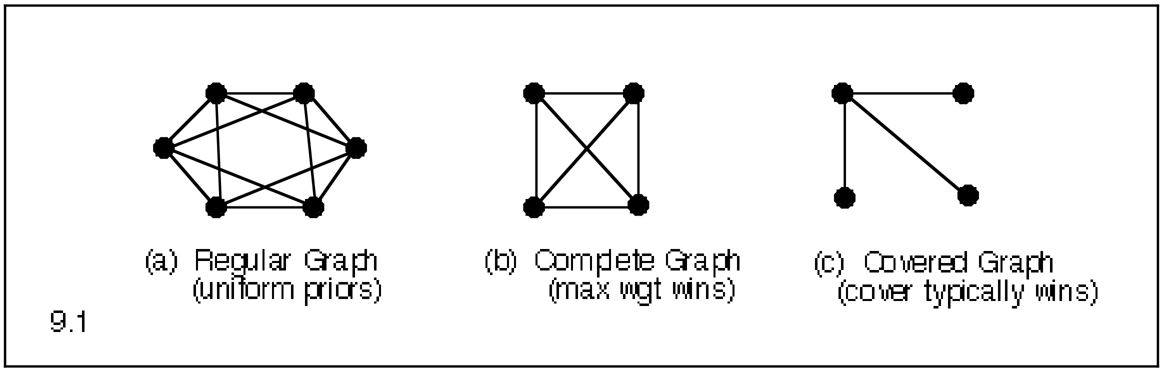

Fig. 9.1: Some anigraf forms can be inferred rather easily from priors.

A second piece of information gleaned from priors will be the ranking of the degree of vertices in Q's anigraf. An example will be given shortly.

But with sufficient memory capacity, P can do much better. He can keep track of the relation between his own choices and Q’s choices for actions, in effect constructing an n x n table of conditional probabilities. Such a table shows the odds for Q’s choices given P’s choice. If P and Q have identical graphical models, the entries will be all 1's along the diagonal, with 0's everywhere else. When P and Q differ, however, the entries will create a more complex pattern. Decoding this pattern allows P to infer Q's internal graphical form -- his mental model for the domain. This is a first step in creating anigrafs who are able to view themselves and others from an overarching perspective –a superego if you will.

9.2 An Example

Let the psychiatrist's model P be a tree of four nodes {A, B, C, D} with C as the a covering vertex (Fig. 9.2, left.) Anigraf Q is known to have the same four goals, namely {A…D}, but Q’s relations between these goals are unknown. P knows that Q's answers will be honest, and that Q also uses a Condorcet tally to calculate winners. The depth of each preference order is 2 (kd=2). Rather than asking Q questions as a typical psychiatrist, P makes note of Q’s actions for sets of input weights on the four mental organisms. This is easy in this context because the same weights also apply to P’s choices. After observing Q’s actions to a host of such "questions", P is able to construct the following table of conditional probabilities:

Table 9.1. Conditional probabilities of Q agreeing with P.

		Q's winner			
		A	B	C	D
P's winner	A	1	0	0	0
	B	0	1	0	0
	C	0.1	0.1	0.8	0
	D	0	0	0	1

The entries give the probability $p(X_q \mid Y_p)$ of Q picking the action $X_q \, \varepsilon \{A, B, C, D\}$, given that P picks Y_p from the same set of actions, for the same set of weights.

P also knows the following:

(i) The priors on his own winners, which for our example are calculated to be (0.04, 0.04, 0.88, 0.04) for alternatives A - D respectively.

(ii) Estimates of the priors for Q, again calculated to be (0.125, 0.125, 0.71, 0.04).

The psychiatrist P can now recover Q's mental model (and vice versa if Q were to have access to P's winners.) The analysis is as follows:

Step1. Conditionals: *Whenever P chooses A, B, or D as the winner, then Q also chooses the same winner* (Table 1.)

At first blush, one might infer that Q's preference orders (i.e. neighbors in Q's anigraf) for A, B and D must be identical to P's neighbors for the same alternatives. But this is not correct. All we know for sure is that SOME of the neighbors for these actions are very likely the same. In other words, if each node or agent in P's model has a smaller number of second choices, then if that alternative wins for P, the same alternative can only have an equal or greater weight for Q. Hence P deduces that Q 's model has at least the following edges: AC, BC, DC, and remains uncertain about the remaining possible edges, namely AB, AD, BD.

Step 2. Priors: *The odds for Q picking D as a winner are very low (0.04) and are the same as for P picking D. The odds for C in both cases are very high, namely 0.7 and 0.8.*

Guiding Hypothesis: The rank ordering of vertex degree in any graphical model is the same as the rank ordering of priors.

Because Q's priors have three levels, namely ~3/4, 1/8, 1/25, and because the maximum degree is 3, P concludes that the vertex for alternative C has degree 3, the vertices for A and B have degree 2, and the vertex for D has degree 1, the latter being the same as P's model for D. Hence Q's model has edge AB, but not edges AD or BD, as shown in Fig. 9.2.

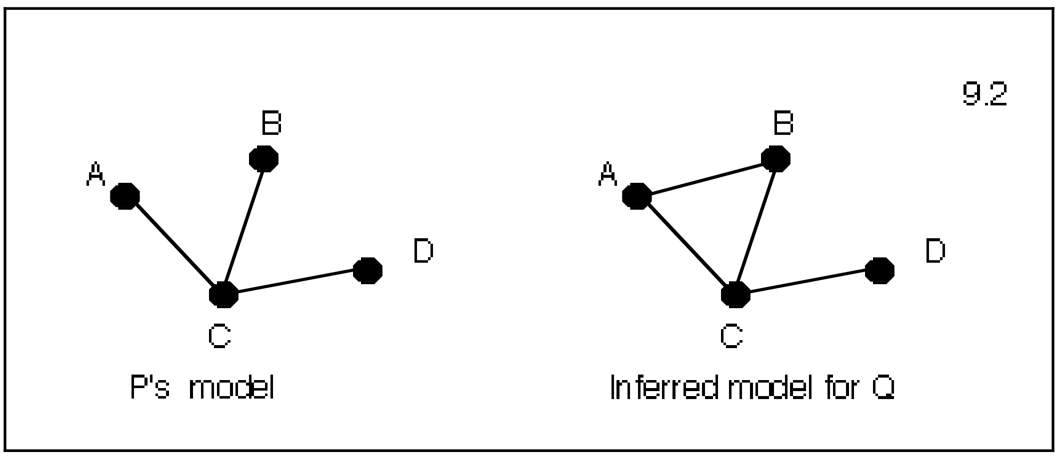

Fig 9.2: P has the internal model shown on the left. Based on Q’s choices for winners, P infers the model at right for Q.

Using the subscripts p, q to identify the anigraf, we note in passing that the probabilities $p(A_q | C_p) = p(B_q | C_p) = \sim 1/12$. This is consistent. Because both A_q and B_q gain strength in a vote from each other as well as from C_q, they will win more often than A_p and B_p. The difference will be absorbed by C_p.

9.3 Poker Intelligence

Most psychiatrists do not reveal their own choice for any given set of "questions". Clearly, without knowledge of P's winners, Q is at a considerable disadvantage, if indeed his answers honestly reveal choices. P knows Q's internal mental model that guides decision-making, but not vice versa. If P and Q are playing for some shared resource, P will eventually dominate, for he can predict all of Q's choices in any situation. Adding new anigrafs R and S to the group will present no new hurdles for P, provided his memory capacity for past winners is sufficient. Thus, P has a strong advantage over others in any competitive game. Many would regard P as the most intelligent of the four because he most quickly grasps the mind-sets of the others.

Recognizing P's advantage, what strategies should competitors Q, R, and S take? Should answers (and consequent actions) be random? Or perhaps misleading answers are better, implying a mental model different from their true preferences. Are these subterfuges worth taking actions that are not really rewarding in the short term? What if circumstances change, and now cooperation is mutually beneficial? Will each player lack the other's trust, or at a minimum be confused by new choices inconsistent with earlier evidence? Why should P come to believe that Q is now giving honest answers? Our social setting is quite different from the case when one of the parties is an automaton, where any mismatching of models can only come from a defect or fault in the machine, not through intentional strategizing. Elsewhere, there has been much effort studying best strategies to restore cooperativity. In the face of non-cooperative behaviors, "tit-for-tat" is one popular choice (Axelrod, 1987.) However, for many situations, honesty is the best policy. But such a policy does not exclude the dominant, more intelligent anigraf from bargaining, or attempting to modify another's mental model to advantage.

9.4 A Group Setting

When anigrafs reflect on cooperative vs. competitive choices, rewards obviously play an important role. Voting behaviors are influenced, or, in anigraf world, the weights placed on the individual actions of the mental organisms will be adjusted. Classical game theory casts these rewards as gains or losses. Here, in contrast with the poker game where each player strives to win as much as possible from others, the gains and losses will apply to the entire population of anigrafs. In other words, the rewards are dispersed throughout the society, according to votes of the anigrafs that comprise this social structure. The optimization problem is to maximize the assignment of rewards to the group members.

In this scenario, just like in all previous anigrafs, we assume that there is one shared global model for the similarity relations between the items to be bartered. Initially, all individual anigraf's models are consistent with this global model. Thus, one anigraf may prefer certain fruits over others, and these over various vegetables. Whereas other anigrafs may prefer vegetables to fruit or meat. But in spite of these different preferences, all have agreed at the outset that fruits are more similar to vegetables than they are to meat. The global similarity relationship among the items of interest establishes the exchange value for any individual anigraf. First we outline how consensus is achieved without any deal-making, and then proceed to a more complex, strategic bartering example.

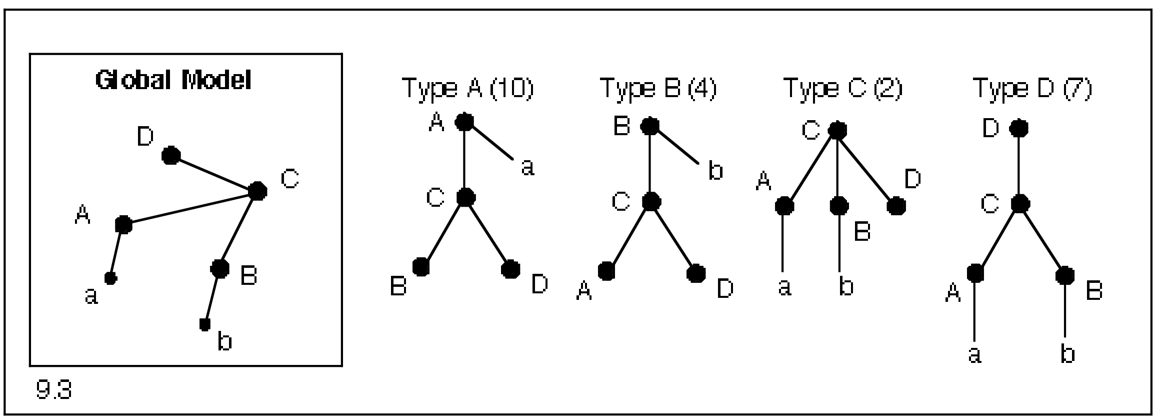

Fig. 9.3: A community is defined by a global model (left). Four subgraphs of this model identify the varieties of anigrafs who belong to the community (right).

Figure 9.3 illustrates using a very simple global model (left). There are four different types of items, or rewards {A, B, C, D}, two of which come in half quantities, designated a,b. For simplicity, all items have positive attributes, and we assume that larger amounts of items are more preferred. Each of these items is associated with one type of anigraf, with that item being their preferred reward. Their graphical forms are shown at the right. Note the consistency with the global model. The number of members of each type are given in the parentheses. (Members of each type are assumed to vote as a block.) Normally, our anigrafs would truncate their preference orders at level two in the graph, with all remaining alternatives being treated indifferently, regardless of the actual depth in the ordering. Here, because of the simplicity of the global model, such truncation only affects the type D anigrafs, with lesser amounts of items A and B, namely a, b, lying at the lowest, fourth level.

With the populations of each type as shown, all anigrafs seek help from others in gathering their own first choice. This process proceeds by having members of all types work together for the society's first choice, then the second choice of the society is gathered, etc. There may not be enough time to complete all the harvesting, so the agreed order becomes important.

The resulting social order S^*_1 of our first Condorcet vote by this population is the following harvesting order: $S^*_1 = (C > A > D > B)$. In other words, given the current anigraf forms and their group sizes, the greatest social benefit is when C is harvested first. Thus, although C is the choice of the smallest sub-population, the tally favors first gathering "fruit C", and then moving on to A's favored reward, then D, with B being the last to be harvested.

9.5 Model Manipulation: Deal-Making

Typically, we would terminate our aggregation process after this vote. But our anigrafs have now reached a level of intelligence that allows them to read the mind-set of others, opening the door to manipulating models. What will happen to the social outcome for harvesting {A…D} if one type of anigraf could convince another type to change its mental model – i.e. to accept a small revision in the graphical relations of Fig. 9.3? Obviously this can succeed only if both types benefit in the social outcome, and if the consequent change in the global model is not inconsistent with the preference orderings of the remaining types of anigrafs.

To illustrate, let each type of anigraf pick as their leader that anigraf with the greatest ability to deduce the graphical relations of others. For example, let B* be the leader of type B anigrafs. Note that in the first vote, B anigrafs lie at the bottom of the social order. But B* observes that their group position as well as that of type D anigrafs can be improved if their leader D* can be convinced that there is now a similarity relation between items B and D. This requires adding the edge BD to the global model. But note that such an addition will not affect the preference orderings for A and C. Hence A and C should have no immediate grounds for objections. B* succeeds in her negotiation with D* and

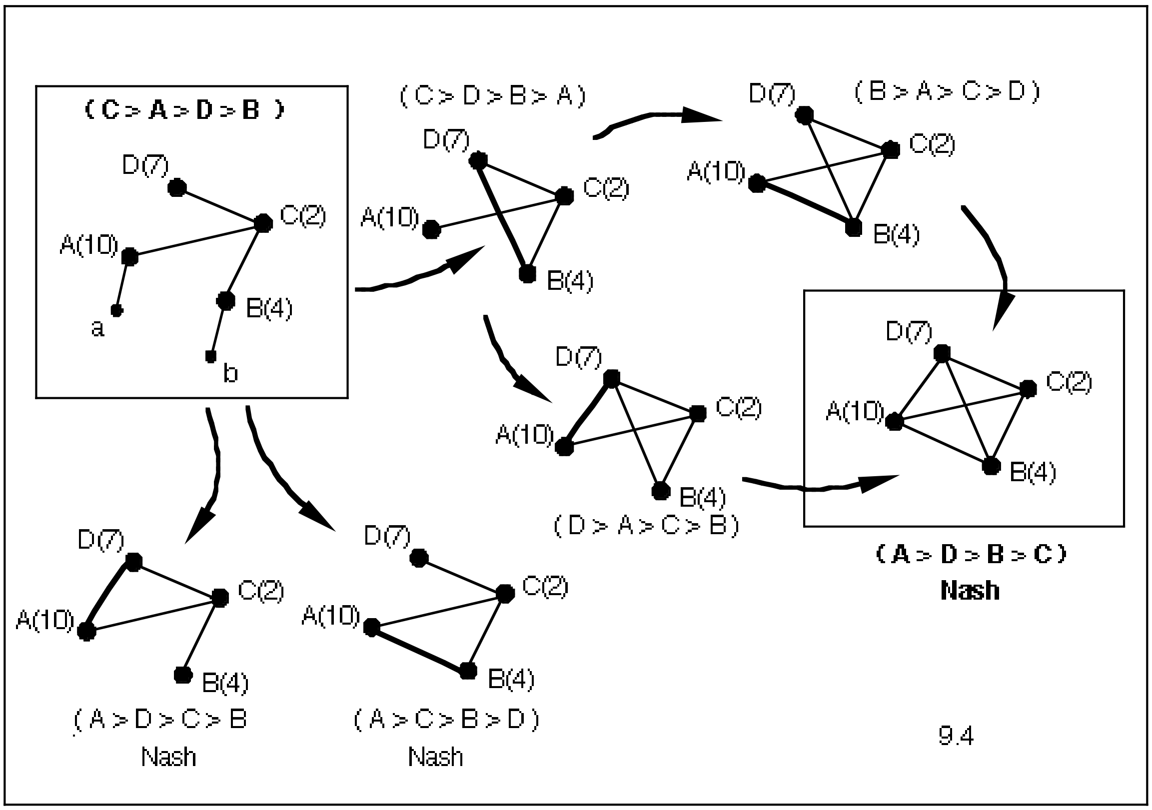

Fig. 9.4: Several sequences showing the effects of different model manipulations on final outcomes.

types B and D anigrafs then modify their models and preference orderings, leaving types A and C unaffected. They now request another the tally. The new social outcome becomes S^*_2 = (C > D > B > A). So both B and D have improved their positions in the ranking at the expense of A! Of course, now A* realizes with the new global model, there is a benefit in establishing a relationship with either B* or D*. If A* and B* reach an agreement using the revised global model, then again the preference orderings of the remaining two anigraf types are unaffected. But now a third tally yields yet another social ordering: S^*_3 = (B > A > C > D). In most cases, if the deal-making continues, then everything will be deemed similar to everything else, and the final global model will be the complete graph K_n. The resultant social order is S^*_{final} = (A > D > B > C), where the ordering is determined by the sizes of the sub-populations.

When edges can only be added but not deleted from a global model, K_n is an obvious Nash equilibrium: no two players can further negotiate to improve their positions. But other fixed points in the social order are also possible, even with the constraint of edge additions only. Consider again the original global model, but with A* and D* negotiating first. Then S^*_2 = (A > D > C > B). This also is a Nash equilibrium. Yet another is shown in Fig. 9.4.

9.6 Evolution and Global Orders

The above are simple examples of the many possible paths in the evolution of a global model for a set of anigraf types. Each step in the evolutionary process improves the position of two groups in the population, without (initially) altering the cognitive structures of the other members. For four alternatives, there are twenty possible anigraf types, and 30 global models, making possible many evolutionary paths even for so few alternatives. Just which sequences will be picked at any instant will rest on the equivalent of a coin flip. Each of these sequences will end in a fixed point that need not be the complete graph K_n (D. Richards, 2003). Hence there will be many sub-optimal fixed points, which are stable until a new element in the social structure is introduced (such as a link to a new population.)

In the biological world, these fixed points are Natural Modes: tight clusters of highly correlated properties. These modes are robust to change because they are associated with very successful designs (Thompson, 1968; Huxley, 1972; McMahon, 1977; Alexander, 1983) Analogously, in the cognitive world of anigrafs, the fixed points are those cognitive structures that are robust predictors and highly successful and stable models for social encounters. Clearly, the Natural Modes of biology and the Cognitive Modes of the mind are at least loosely coupled, because designs of one domain influence designs in the other. Although predicting the dynamics of the evolution of modes seems near impossible, the structure of the fixed points may be accessible. In the cognitive arena of anigrafs, this structure will be the ubiquity of particular graphical forms and their relationships. Metagrafs begins this journey.