Concepts from Pythagoras

James Bedford-Roberts
Information Management Laboratory
HP Laboratories Bristol
HPL-91-22
March, 1991

policy, right, role, responsibility, organization

The Pythagoras project aims to enable the explicit representation and manipulation of rights, roles and responsibilities by information users within an organization.

In this document we develop an interrelated set of concepts concerning rights, roles and responsibilities. The purpose of these concepts is to act as a catalyst in the formation of a common and clear terminology, useful both between different people and between people and machines.

Internal Accession Date Only

1 Introduction

1.1 Function and form

The Pythagoras Project at Hewlett-Packard Laboratories, Bristol, England aims to enable the explicit representation and manipulation of rights, roles and responsibilities (loosely collected as policy) by information users within an organisation. This aim reflects our belief that there is a new field emerging in Computer Science concerned with managing policies within human organisations. In this document we develop a provisional model for policy management.

This document is an offshoot from discussions about policy. Perhaps it is best to read it as a personal record of those discussions. Like so many discussions, it is incomplete and contains diversions that have varying degrees of relevance to the whole. Furthermore it does not set off to establish any single point but rather to map out a domain of interest.

On their own the concepts from the model form a terminology useful for the exchange of ideas between computer scientists working in this area. Our hope is that our provisional model will prompt the evolution of a standard terminology in the field. In addition our model can be used to provide clear terminology for the interface between a computer that helps manage policies and its users.

The model is presented through a set of concepts. In section 2 we describe some primitive concepts. These concepts provide the foundations for subsequent ideas. In sections 3, 4 and 5 we introduce more complex and more useful concepts, often defined in terms of the more primitive concepts from preceding sections.

When a concept is defined, it is printed in bold type face. If the concept definition builds on other concepts, these other concepts are printed in italics. Italics are also used when a term appears before it is defined and sometimes simply to give emphasis. There is a glossary in section 7 on page 20.

1.2 Background

People who are part of a human organisation (such as a factory, office or hospital) behave distinctively compared to people who are not. There will be times when a person's behaviour is constrained by an organisation to which they belong. Such constraints are often conveyed as policies (for example, Bob must receive permission from Ann before going on holiday). We are interested in developing a model that can describe the way in which organisations manage people through policies.

Our model can be used to explain the observable characteristic behaviour of human organisations. Characteristic behaviour is, if you like, the tip of the iceberg; most of what goes on within an organisation is below the observable surface. In particular most of the activity goes on within the minds of the individuals in the organisation. We are not concerned to understand what might really be going on within people's

minds, just so long as we can devise a mechanism that can explain observable behaviour. Over all, we are interested in behaviour that is observable (else it is not worth computer support) and can be described (else computers cannot support it).

2 Basic concepts

2.1 Entities

We start our model from the primitive definition of an *entity* as a "thing". An **entity** is an atomic unit, that is, something we can view as a whole. At another time it may be possible to view the same entity as being a composition of other entities. All entities have *state*. **State** is the condition in which an *entity* is, its mode of existence as determined by circumstances. The state of an entity reflects the collective states of all the entities of which it is composed.

2.2 Observers and the system

The process of creating a model of a human organisation involves observation, this is carried out by an observer. The observer is a user of our model and *not* a part of the model.

The **system** is an *entity* conceived by an observer to delimit the extent of what he wishes to observe. The system is not just defined by physical boundaries, it may also be defined by temporal boundaries and by root causes or atomic (undecomposable) entities. All these different boundary types may be interdependent, for example a system's physical boundary may depend on the time. In short, at any point an observer may stop any progressive analysis by saying that he has reached the system boundary.

Observers may observe many (possibly overlapping) systems at once, but any observation they make must be associated with a specific system. The system *is* part of our model.

2.3 Reality and perception

We define an **individual** as an *entity* capable of having perceptions. By 'individual' we usually mean a person: a human being (but see section 4.2 on page 12). In this section we emphasise perceptions as being sensations of the outside world as it currently is, but perceptions could equally be memories of past worlds or imaginings of future worlds. All these perceptions give rise to an individual's beliefs, which in turn effect the individual's behaviour.

Ultimately, if we are to be able to model human behaviour within organisations then we must capture more than just the observer's reality; we must capture people's perceptions as well.

In fact, "reality" (as perceived by an observer) has no direct effect on the behaviour of others. No matter how rational they are, they will be responding to their perceptions

rather than to "reality". In fact we do not rely upon there being any independent reality, just perception, and similarly it does not matter if there is no such thing as knowledge, just so long as there are beliefs. For us reality and knowledge can be just labels that people give to reflect their faith in *their own* perceptions and beliefs. We can talk about reality and knowledge in our model because this model is intended to reflect the *perceptions* of an observer. When we speak of things as if they were real we mean they are perceived by the observer, when we apply our model to make some factual statement, we mean that the statement is the belief of the observer.

This stance allows us to avoid some difficult philosophical problems about how we can know which entities have perceptions and which do not. After all, only the entity itself can know whether it can perceive. Our line is that if the observer chooses to believe that an entity within the system is an individual then we are entitled to state it as a fact. Our model is only concerned with the perceptions of the observer.

It is unclear what impact, perception and belief should have on the support offered by an organisational support system. In particular we are not advocating that the system intentionally support deception and dishonesty. Nonetheless it does suggest that if the system is to be able to query knowledge supplied by different people, then it should be able to handle contradictions and undefinedness as well as the more straight forward truth and falsity.

2.4 State changes

In the previous section we suggest that individuals respond to their perceptions, and *not* directly to the observer's "reality". Thus in this section about system state changes we are more concerned with perceived state changes than with "real" state changes. These perceptions may be sensations from the outside world as it changes, or visions of how the world could be, or memories of changes that have taken place. For example, both the perception of 'being paid tomorrow' and the perception formed when entering a bath can influence behaviour. It does not matter if the perception comes from a day dream, a memory or (the observer's) reality. We introduce the term 'effect' to cover all these perceptions.

An **effect** is a perception of a state change.

There are different types of effect. A **goal** is an *effect* desired by an *individual*. An **ability** is an *effect* accomplishable by an *individual*. An **action** is an *effect* accomplished by an *individual*. An **event** is an *effect* that has occurred (according to the observer)[1].

We can present all the concepts of 'goal', 'ability', 'action' and 'event' as kinds of effect with a special distinguishing attribute. We show this in Figure 1. This arrangement

1. Type theorists may prefer to think of effects as types and events as instances. We could have defined events more directly as state changes (observed by the observer), but we prefer to see them as kinds of effects.

allows us to make observations once for effects rather than repeatedly for each of goals, abilities, actions and events.

Type of Effect	Associated Attribute
Goal	Desired
Ability	Accomplishable
Action	Accomplished
Event	Occurred

Figure 1 Attributes of different types of effect

Some attributes are related. For example, we may consider an action as event perpetrated by an individual; so any effect with the attribute 'accomplished' also has the attribute 'occurred'. Similarly, at the point when an effect becomes 'accomplished', it must hold the attribute 'accomplishable'.

2.5 Composing and decomposing effects

Effects can be composed and decomposed[1]. It is useful to introduce terms to refer to the products of composition and decomposition of an effect. A **super-effect** is an *effect* resulting from the composition of other *effects*. A **sub-effect** is an *effect* resulting from the decomposition of another *effect*.

In section 2.2 on page 3, we mentioned that the observer may use the system boundary to limit progressive analysis of the system. In particular the observer can set up the system boundary so as to limit progressive composition and decomposition. Thus an **atomic-effect** is an *effect* considered to have no *sub-effects*. Similarly, a **root-effect** is an *effect* considered to have no *super-effects*. The observer's choice of a set of atomic and root effects is arbitrary, effects are *not* intrinsically atomic or root. Specifying the atomic and root effects of a system is a part of specifying the system boundary.

Without further definition we will apply the prefixes 'super', 'sub', 'atomic' and 'root' to entities, effects, goals, abilities, actions and events.

1. Composition and decomposition relationships are based upon some notion of equivalence between the super-effect and associated sub-effects. The usual sense of equivalence for effects means that the realisation of the sub-effects is equivalent to the realisation of the super-effect. Equivalence for entities is based on similarity of physical make-up.

2.6 Deriving the attributes of effects

For some attributes, any composition of effects sharing that attribute will make up a new super-effect that also has that attribute. In these cases the super-effect can be seen to have derived its attribute from its sub-effects. Ultimately the derivation can be traced back to some combination of atomic effects also displaying that attribute. For example, an effect may be deemed accomplishable by virtue of there being a decomposition into entirely accomplishable atomic effects. The same applies for the attributes 'accomplished' and 'occurred'. The 'desired' attribute of goals is an exception to this.

The desirability of a goal derives from its super-goals rather than its sub-goals. For example, Mick's desire to take the train to Bristol (goal) is derived from his desire to visit friends living there (super-goal). It is not derived from his desire to buy a train ticket (sub-goal).

Ultimately all desire can be considered to stem from root goals and all ability from atomic abilities. The flow of desire and ability from these sources is controlled consciously by individuals through the process of planning.

2.7 The planning process

Planning is a well studied area of AI [Charniak 85]. We mention it here only to indicate its importance in linking perceptions and behaviour, or more specifically, in linking goals and actions.

When an individual first appears in the system (it does not matter how), he has a set of root goals and a set of atomic abilities. The individual's thinking is driven by the desire to achieve his root goals. It is likely that an individual will find that most of his root goals are not accomplishable from his atomic abilities. To make each of his goals accomplishable, the individual can use two tactics:

- Decompose an existing goal into a new set of sub-goals that are individually easier to accomplish. For a given goal, an individual may be able to think of many different sets of sub-goals. It may be hard to decide which set of sub-goals to pursue.
- Compose a new ability from existing abilities, with the objective of creating an ability that would achieve a previously unaccomplishable goal.

This mode of working whereby individuals break down goals and build up abilities is shown in Figure 2.

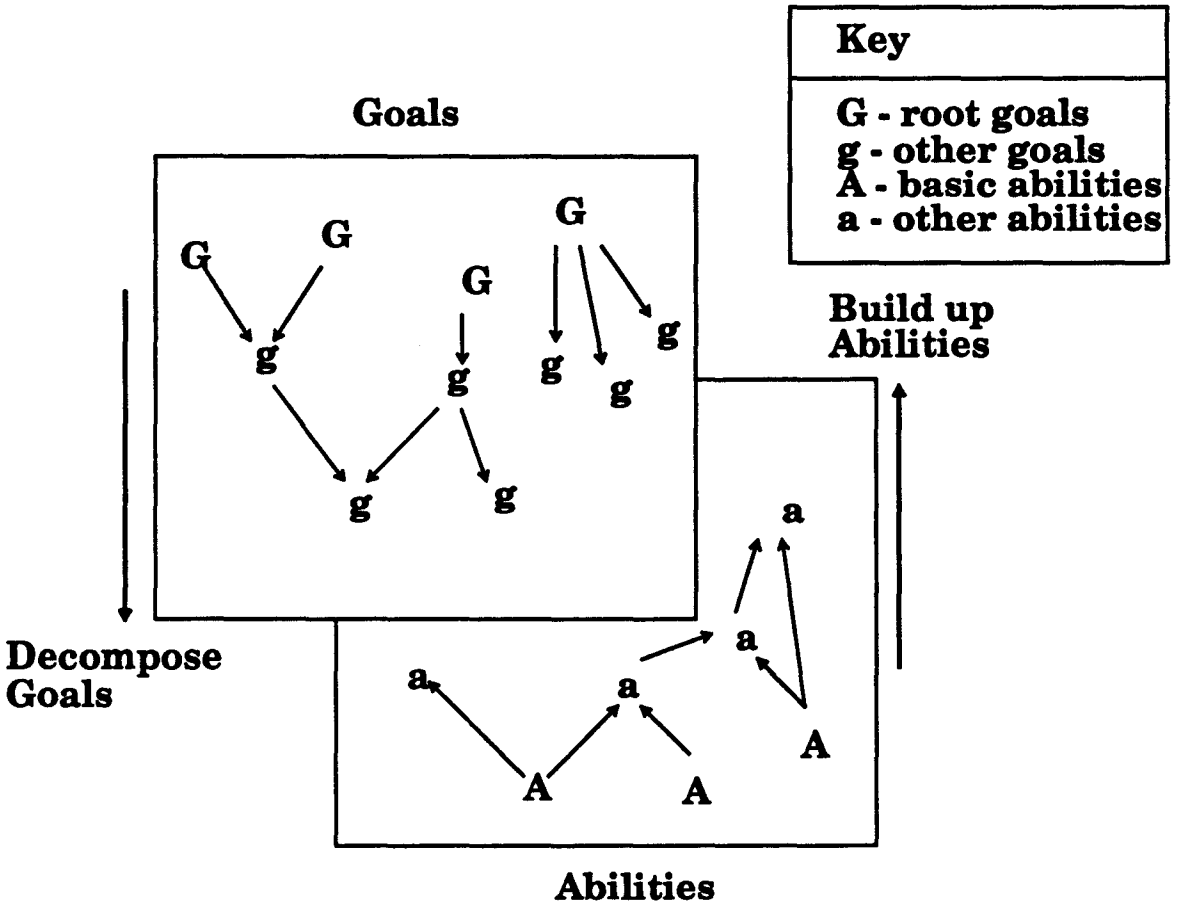

Figure 2. Breaking down goals, building up abilities

2.8 Satiable and insatiable goals

In the previous section we refer to assembling abilities to accomplish goals, but insatiable goals can never be fully accomplished irrespective of ability.

Some goals are satiable others are not; for example earning one hundred pounds is satiable, but being as rich as possible is not. An **insatiable goal** is a *goal* with no associated condition of accomplishment.

It is not always clear whether a goal is satiable or not. The goal of eating is satiable in the short term but insatiable in the long term, that is, a large meal may satiate for a while, but no meal suppresses appetite permanently (unless it is lethal). Satiability is not an intrinsic quality of goals, Tristan may have an insatiable desire to play rugby for ten years and then suddenly get fed up with the game. Alternatively, you could

assert that playing rugby was always satiable for Tristan and that it just took a long time for him to satiate his desire.

Note that direct removal of the desire for a goal does not count as satiating it; the desire must be removed through accomplishment. This is a rather grey area because it is not always clear what the reason for the loss of desire is. For example, Tristan may stop playing rugby because he is fed up with the game (satiated goal) or simply because he got injured (unsatiated goal).

Other insatiable goals could simply be concerned with maintenance of some state of affairs such as 'keep the juggling balls in the air' or 'keep the department running smoothly'. From the "system state" point of view we could associate satiable goals with state transitions, and insatiable goals with state maintenance. The match is not perfect though, for example 'getting richer and richer' can be viewed as an insatiable goal involving state transitions from one level of wealth to another. The mismatch reflects the question of how we describe the system state, and consequently what constitutes maintenance or change of that state.

2.9 Planning by numbers

Much of planning concerns choosing one plan over another. The decision will often be made quantitatively. How desirable is the goal? How much effort is required to accomplish it using plan A? Would less effort be required using plan B?

The quantitative nature of this process suggests that perhaps plans could be computed. Conceivably root goals and atomic abilities could be assigned numeric values to indicate their respective desirability and accomplishability. A computer may then deduce further values for sub-goals and super-abilities using rules of composition and decomposition. Ultimately all plans may be assigned numeric values for cost (required effort) versus benefit (gratified desire).

The problem is that people do not actually plan in this way. They perform much more intuitively. Intuitive decisions are things that people do well. It is unlikely that people would ever be happy to give up control of this area of their activity. Thus the real point of this section is to observe that there are some aspects of organisational behaviour that are best left in the hands of people.

3 Working together

3.1 Different abilities

Different individuals have different abilities. In some cases the goals of one individual can only be accomplished using the abilities of another. Alternatively, the goal-holder might be able to accomplish his goal, but perhaps not as successfully, or quickly, or easily as another individual could. Or perhaps it is just that the goal-holder does not have sufficient time.

In all these cases it might be better for the more able individual to perform the goal on behalf of the individual that desires it. This seldom happens because individuals only act according to their own goals. The problem is: how can goals be transferred between individuals for the sake of the greater common good? We approach this problem by considering the concepts of *trust* and *commitment*.

3.2 Trust and Commitment

Trust is the belief of an *individual* in the integrity of another. **Commitment** is the intention of an *individual* to act with integrity towards another. Trust and commitment are often specific to the pursuit of a particular goal (for example, Sally might trust Ian to set the burglar alarm, but she might not trust him to reverse the car into the garage).

We do not address the question of how different individuals come to feel trust or commitment. We assume that trust and commitment are inherent yet transient characteristics of individuals. One individual may feel trust or commitment towards another, without the trusted or committed-to individual being aware.

3.3 Rights and Responsibilities

Trust and commitment are not useful on their own, but can signal the correct conditions for *delegation* (of a goal). **Delegation** is the arrangement whereby one *individual* pursues a goal on behalf of another *individual*. Delegation requires a specific goal for which one individual extends trust and the other commitment. If the individuals concerned recognize these conditions, they may negotiate a *cooperation*. If the negotiations are successful then the cooperation is established and as part of this the goal is delegated. When this happens the trusting individual relinquishes control of a goal and the committing individual promises effort to accomplish that goal.

Somewhat arbitrarily we suggest that the delegation is established once the trusting individual has received a *right* from the committing individual, and the committing individual has received a *responsibility* from the trusting individual. This exchange of rights and responsibilities constitutes a binding between the cooperating individuals. A **responsibility** marks a goal held by an individual that has been delegated *from* another. A **right** marks a goal held by an individual that been delegated *to* another.

A difficulty some people find with these definitions of right and responsibility is that they relate too strongly to delegation[1]. Indeed for many real examples of rights and responsibilities the delegation process is remote from the responsible and right holding parties. For example, some citizens are born with civil rights, which they obtained without participating in any delegation process. In this case the delegation process occurred between their ancestors and their ancestors' government. (Their ancestors

1. More naturally, we would expect people to suggest that rights and responsibilities are both rules. They might suggest that rights work for you whereas responsibilities work against you.

wanted civil liberties for themselves and for their descendents, and managed to delegate this goal to their government, which subsequently took responsibility in the matter.)

So, we accept that the delegation process may be remote from the resulting rights and responsibilities. In spite of this, identifying the delegation associated with rights and responsibility is important. This is because the delegation process provides the context from which the strength of the underlying trust and commitment can be measured.

In the first instance it may be easier to identify rights and responsibilities as transferred goals than as by-products of a delegation process. Once a pair of rights and responsibilities have been identified it may then be profitable to search for the delegation process by which they were established.

In organisations, rights and reproducibilities often correspond to specific trusts and commitments between individuals and representatives of the organisation. For example, a manager may trust an employee to be punctual for a meeting. This may become a responsibility of the employee and a right of the employer.

3.4 Client and Server

We will call the *responsible individual* the **server** and the *right*-holding *individual* the **client**. At any stage the server may withdraw the right he has extended to the client and the client may withdraw the responsibility he has extended to the server. If either party withdraws their rights or responsibilities it signals the end of the delegation. The withdrawal may be initiated unilaterally or it may follow some negotiated procedure.

The concepts of trust, commitment, rights and responsibilities are summarized in the Figure 3.

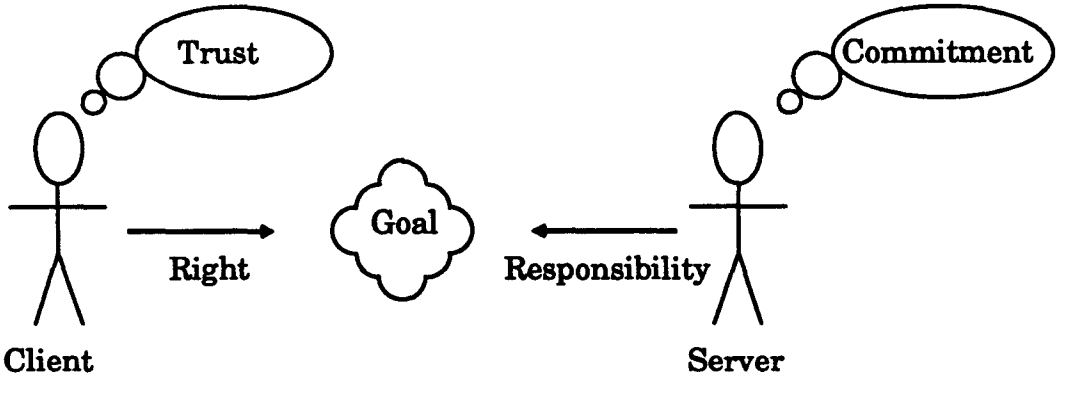

Figure 3 Attributes of clients and servers.

3.5 Cooperation

It is not enough that two parties are both aware of their mutual trust and commitment concerning a particular goal. If the goal is to be delegated it must be part of an arrangement that is beneficial to both individuals. Ultimately the only way that individuals can benefit each other is through acting as servers in a delegation. We define a **cooperation** as a set of *delegations* balanced so as to be beneficial to all the *individuals* concerned. The negotiation that follows the recognition of potential delegations and their establishment concerns trying to group delegations into a cooperation. A cooperation may involve many individuals and many delegations. Normally each individual acts both as a client and as a server in different delegations from the same cooperation.

Cooperations are fragile; if any delegation ends then the balance may be upset and all the remaining delegations may end soon after. Large cooperations are particularly fragile, but they can also be particularly beneficial to the individuals concerned.

Cooperations are often open-ended. Rather than describing the delegation of specific goals they may establish long term trusts and commitments that ease the formation of future delegations. Mutual benefit can be achieved directly through return of service, but perhaps more often it is achieved indirectly through the payment of money. This is the usual form in which organisations establish cooperation with employees.

4 Enterprises

4.1 What is an Enterprise?

An **enterprise** is a group of *individuals cooperating* to pursue a set of common *goals*, or using the definition of a cooperation we can say that an enterprise is a cooperation based on a set of common goals. There need not be any direct relationship between the goals of the enterprise and the personal (undelegated) goals of any individual within the enterprise. There need not be any relationship between the goals of the different individuals in the enterprise; they may share goals not associated with the enterprise and they may even hold conflicting goals. The goals of an enterprise may be club-like, inward looking, concerning the mutual benefit of the members or they may be expedition-like, outward looking, concerning some form of conquest over the outside world.

The individuals within an enterprise are delegated goals from the enterprise and affect actions in order to achieve them. An individual takes part in an enterprise because it is to his advantage. This may either be because the goals of the enterprise coincide with the goals of the individual, or because the enterprise is able to offer the individual some incentive for joining.

4.2 An enterprise is an individual

An enterprise is itself an individual. This follows from our definition of an individual as an active entity motivated by goals; any collection of active entities motivated by goals can itself be seen as an active entity motivated by goals. An enterprise may be more than the sum of the individuals within it. For example, an enterprise such as a newspaper can have its own funds, not associated with any individual, and can be sued for liability independently of any of its members. Because enterprises have no intrinsic physical form, whenever we refer to an enterprise performing some physical action (such as negotiation) we mean that a *member* of the enterprise performs the action on behalf of the enterprise.

Every enterprise will have a constitution by which the members animate their enterprise. The constitution could take any form, it might be democratic on every decision, or perhaps different members have different areas of responsibility, or perhaps the enterprise is run autocratically by a single member.

The enterprise concept is important because it allows different people to share authority and ownership (see sections 5.2 and 5.3). At the same time, enterprises introduce ambiguity because none of the enterprise members may be entirely accountable for the actions of the enterprise. They also introduce ambiguity when it is unclear whether a person is acting on their own behalf or on the behalf of the enterprise. Both kinds of ambiguity can be tackled by making clearer the relationship between people and roles, we mention this issue in section 5.1.4.

4.3 Combining enterprises

Given that enterprises are individuals and that enterprises contain individuals, it follows that enterprises may be nested one inside another. Note that there is no reason why individuals may not be members of more than one enterprise at the same time.

All cooperations may be recognized as enterprises with the goal of mutual benefit. Cooperations may also cut across enterprises. In our general model cooperation between enterprises follows exactly the same pattern as cooperation within an enterprise, and for that matter cooperation between individuals.

4.4 Recruitment

Since an enterprise can not perform its own actions it must trust its members and achieve its goals by establishing delegations. Business organisations tend to form hierarchically arranged enterprises whereby the primary set of goals associated with an enterprise are progressively delegated and decomposed from the original enterprise to member enterprises and eventually to single individuals. This can result in the delegation of extremely abstract goals to members near the top of the hierarchy. Commonly, a member may accept the goal of accepting further unspecified responsibilities from another member of the enterprise. Such non-specific commitment is a typical requirement for military recruits. Open-ended arrangements such as assuming a re-

sponsibility to assume further responsibilities are also common; the member is usually protected by having the right to leave the enterprise if he wishes. The underlying point here is that although the model is fairly basic, it can nonetheless be used to describe complex situations.

The recruitment of individuals into an enterprise can occur in many ways. Recruitment is a negotiation process between a potential client-server pair, where the client individual happens to be an enterprise. Being a member of an enterprise is exactly equivalent to being a server in a delegation where the client individual is an enterprise. The new member assumes a set of responsibilities (and possibly rights) to the enterprise

5 Other concepts

In this section we introduce terms that are commonly used in discussing the structure of organisations. We explain these terms using the more primitive terms defined in the previous sections.

5.1 Procedures, strategies, roles and cooperative arrangements

Procedures, strategies, roles and policies are commonly used terms in the description of organisations, which turn out to be related by the form of their definitions. We find that they all concern collections of tasks. These collections may be associated with individuals (as in roles and procedures) or with goals (as in policies and strategies). Furthermore the tasks may be expressed as abilities (as in strategies and procedures) or as delegated goals (as in policies and roles). The definitions are summarized in Figure 4, which can be read, for example, as a 'procedure' relates an 'individual' to a set of 'abilities'. The definitions are explained in more detail in sections 5.1.1 through 5.1.4.

	Abilities	Rights/Resps
Individual	Procedure	Role
Goal	Strategy	Cooperative arrangement

Figure 4 Procedures, strategies, roles and policies.

5.1.1 Strategy

A **strategy** is an organized arrangement of *abilities* associated with a *goal*, such that the *goal* can be accomplished by the orderly accomplishment of the *abilities*. Any subtree selected from the directed acyclic graph of goals may form the basis of a strategy. All the abilities of a strategy are associated with sub-goals of the strategy's goal[1].

The actions within a strategy may be organized using conditional constraints such as *if-then*. For example, *if* the car is low on fuel, *then* stop at a petrol station. There may also be sequential constraints. For example, with a goal of getting dressed, putting on your shoes should come before tying up your shoe-laces.

Here, we are most concerned with strategies that are *in effect*. That is, there must be an individual with the associated abilities that desires the goal. Once a strategy is put into effect the goal and its abilities are related, thus the goal becomes accomplishable and the abilities become desirable.

We shall call hypothetical strategies for which no particular desiring or accomplishing individuals are envisaged: 'strategy statements'. We make this distinction in the belief that real organisational systems will handle active strategies differently from strategy statements.

5.1.2 Procedure

A **procedure** is an organized arrangement of *abilities* associated with an *individual*. Comparing procedure and strategy, we note that a strategy may relate different individuals to a common goal, whereas a procedure may relate different goals to a common individual.

Both goals and individuals are subject to decomposition and thus we may identify a sub-individual whose procedure comprises abilities motivated by a single goal, or a sub-goal whose strategy comprises abilities assigned to a single individual.

We take procedure to be synonymous with the term process.

5.1.3 Cooperative arrangement

A **cooperative arrangement** is a collection of *rights* and *responsibilities* associated with a specific *goal*. For example, there may be a policy associated with the goal of keeping a house clean. It could involve visitors being responsible for wiping their feet before coming in and having the right to demand the bathroom be clean and bright before they have to use it.

A cooperative arrangement is like a cooperation except that it does not identify the individuals involved. (Admittedly, the two terms do not seem to suggest this distinction on their own.)1

In contrast to strategies, the decomposition of the goal associated with the cooperative arrangement stops when the sub-goals may be delegated to specific individuals. The amount of decomposition depends on the individuals being considered. For a cooperative arrangement to be "in effect" there must be a real individual desiring the cooperative arrangement's goal and there must be individuals assuming the associated rights and responsibilities.

---

1. Perhaps strategy is not exactly the right word for relating goals to *actions*. After all we might call the single level decomposition of a goal into sub-goals a strategy even if no accomplishing actions were identified.

5.1.4 Roles

We define a **role** as a collection of *rights* and *responsibilities* associated with a single *individual*[1]. The role is defined by the named collection of rights and responsibilities and not by the individual, which may change from time to time, or even be unassigned. An individual may hold more than one role at a time.

An individual is not tied to a particular set of roles, his roles may change as he joins and leaves different enterprises. Furthermore an individual may have a set of roles for each enterprise of which he is a member. Although an individual may hold many roles simultaneously, any action should be associated with a single role.

As in section 5.1.1, we are more interested in roles that exist than in hypothetical roles. For a role to exist we suggest that it must be associated with existing rights and responsibilities. In other words, we are interested in role instances rather than role types. To emphasize our concern with role instances we will prefer the term *hat* to role. This term alludes to the metaphor of people wearing hats. We define a **hat** as a collection of *rights* and *responsibilities* associated with a single *individual*. To reinforce the metaphor we will prefer the term *head* to individual. A **head** is anything that can assume (wear) the rights and responsibilities associated with a *hat*, it may be a person or a group of people with a constitution to bind them. Alternatively, borrowing our definition of an individual, a head is an entity capable of having perceptions.

5.2 Ownership and possession

Ownership is a perceived relationship between an individual (the **owner**) and an entity (the **property**), whereby the owner is believed to have ultimate control of the property. (Entities include things such as information and ideas.)

What is meant by ultimate control? Ultimate control means that an owner has the right to control all other rights over his property. In practice it means that no other individual can use another's property without permission.

Ownership is not real, it is simply a perception. Notice that ownership can be disputed or simply violated. Usually this comes as no surprise to owners; thefts and break-ins are common place. The main significance of ownership is as a compact way of describing a set of behaviours. Ownership confers an individual with a set of expectations about what he and others may do. An individual with a sense of ownership over some property will consider violation of his ownership as a significant event, which may trigger some response. Without that sense of ownership the same event may go unno-

1. Rights and responsibilities that appear as part of role definitions are particularly hard to see in terms of delegated goals as proposed in section 3.3. This is because the delegation process is normally performed by the individual that originally designed the role and not by the individuals that subsequently assume the role. This makes the delegation process particularly remote from the right-holding and responsible parties. Indeed so many rights and responsibilities within an organisation are associated with roles that an alternative definition of rights and responsibilities might describe them as elements of a job description.

ticed. This point is to emphasize that ownership is a state of mind held by an individual.

Ownership tends to be obeyed. Property is frequently protected in proportion to its worth, such that violation is not worthwhile (unless the violator is much stronger than the owner). Additionally, potential violators usually have their own property and prefer to establish violation as a taboo by not violating other's ownership themselves. Finally, respect for ownership may often be supported by the organisation as a whole. In this case it becomes harder for thieves to avoid retribution, even when the rightful owner has little power of his own.

The strong effect of ownership tends to explain its significance within an organisation. Note that in our model it is *not* a fundamental concept from which others are defined.

5.2.1 Ultimate control

Conflict of ownership can arise in various ways. The most obvious is through explicit disagreement, but there are other more subtle forms that arise from the difficulty of making plain what we mean by 'ultimate control', for example:

- through sharing. Anna and Beryl may agree to share X. What mechanism for sharing do they use? If Anna and Beryl successfully agree on the mechanism, then there is no problem, otherwise there is a potential conflict concerning their respective rights of ownership.
- through borrowing. If Anna borrows some property X belonging to Beryl, what rights does Anna have? For example, Beryl may be committed to give Anna ultimate control for a specified period, in this case does Anna become the owner for that period, or does Anna's commitment to return X mean that ultimate control remains with Beryl? Note that borrowing can be seen as a form of sharing.
- through indirect ownership. If Anna owns some property X and yet Anna is committed to obeying Beryl in all matters relating to the property, does Anna really own X or does Beryl? Beryl's control is based on the sincerity of Anna's intention to obey. If Anna is committed in her intention to obey she will probably believe Beryl to be the owner. Nonetheless, Anna is an autonomous individual, this autonomy allows her to revoke her commitments at any point. This would seem to suggest that ultimate control is with Anna.

5.2.2 Ownership and possession

This indirect ownership scenario highlights the difference between ownership and possession. Ownership is about rights, and possession about power. Often possession is related to physical proximity. When this is the case possession is unambiguous because things can only be in one place at a time.

Given that ownership is intrinsically ambiguous (both because it is a belief and because it is unclear what constitutes ultimate control) it is perhaps better to focus on

possession as a less ambiguous cousin of ownership. This would seem to reflect legal processes where "possession is nine tenths of the law". This is not to say that possession amounts to ownership, but rather that in trying to establish the owner of some property, a good first step is to contact the possessor. Where there is no dispute the possessor will either be the owner or he will be related through some commitment to the owner.

Possession does not make all our problems disappear. To start with, a physical entity need not be possessed by anyone (for example, an undiscovered desert island). Furthermore, it is arguable whether non-physical things are possessions at all; for example, some might argue that a piece of paper bearing the representation of an idea is a possession, but that the idea itself is not. Nonetheless, such non-physical things can be owned: this is what patent law is about. In a computer system, one possibility is to start from consideration of the possession of the representation of non-physical entities and work from there.

5.3 Accountability and authority

People are the ultimate sink of accountability because they are the ultimate source of authority. This dualism matches the split between heads and hats.

A head may own many different hats, but he can only wear one at a time. The hat that a head wears defines his rights and responsibilities. These rights and responsibilities define the authority and goals of the head, which in turn lead the head to invoke actions. The head is deemed accountable for the actions it performs while wearing a hat (see).

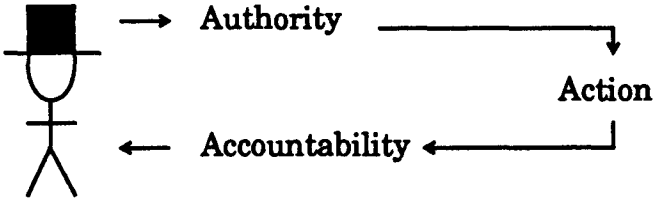

Figure 5 Authority and accountability

Accountability is a difficult issue, perhaps we just mean answerable. The point is that the head can always claim that he was just following the goals associated with the hat. The hat designer and the hat delegator are also to some degree accountable. Passing the buck, makes accountability hard to establish. There is a further difficulty in deciding in what order the potentially accountable parties should be approached. In

this case it seems reasonable to address the perpertrator of the action first. This is all that is intended by deeming the hat wearer accountable.[1]

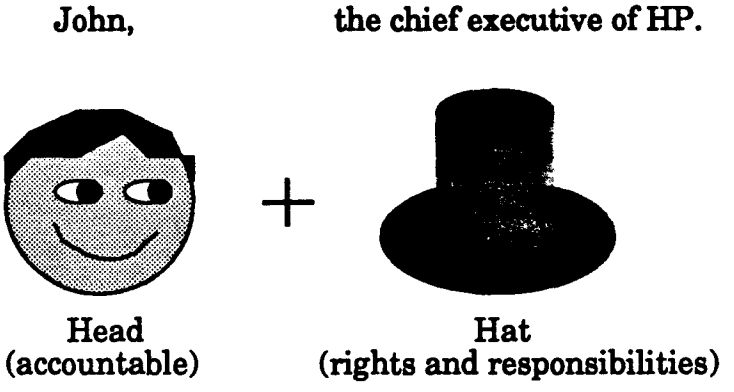

Figure 6 Head and hat

Jonathan Moffett [Moffett 90] has suggested that power and rights are orthogonal components of authority. Taking this view, we propose that power is associated with the head, and rights with the hat. Correspondingly, actions performed by a head whilst wearing a hat can be said to have authority. In our system, by insisting that heads cannot perform actions without wearing a hat we ensure that all actions have authority and not just power behind them. We nonetheless need to be able to accept that in the real world actions are often taken without right.

Actions are thus associated with a head-hat pair such as 'Ann the secretary of the AIM department' or 'John the chief executive of HP'. Some heads and hats are so closely associated that it is hard to name the hat distinctly from the head. For example, when we talk about the head-hat pair 'Jim the engineer' it is hard to name Jim's engineer hat without referring to Jim. Another point of confusion arises when people are just being themselves. In this case it is tempting to think of that person as a head not wearing any hat. In fact people are always wearing some hat. If nothing else they are wearing the hat that represents their personal goals, moral code and civil liberties.

There are philosophical issues concerning whether people can wear more than one hat at once, and what it means when one hat subsumes the rights and responsibilities of another hat. We do not cover these here.

---

1. It is interesting to compare the process of establishing the ultimately accountable party with the process of establishing the owner of some property. To establish accountability for an action you might start with the perpertrator, to establish ownership for an entity you might start with the possessor.

5.4 Policy

There are two types of occasion when it is normal to apply policy. Firstly, when intention and consequent action are separated in time. Secondly, when intention and consequent action are associated with distinct people, or to be more exact, distinct head/hat combinations. In general policy is used to help people convert their intentions into actions when the two have been separated. This could be seen as the most basic problem addressed by Pythagoras, and justifies the project's formal title of 'Policy Management'. By way of a complete definition, a **policy** is a *right* or *responsibility* declared so as to prompt *action* conformant with an intention.

Policy may appear in many forms. A policy may apply to a single specific individual, but more usually it will apply to a group of people. This group is often defined in terms of some common role (for example, it may refer to the role 'engineer' as in the policy: 'all engineers should learn German'), but the group could be defined in any way. A policy may result in a single action, but more usually it will trigger many actions. A policy may be satiated at some point, or perhaps it will apply indefinitely. A policy could even be made and then served by the same head/hat combination, but more often we are concerned with policies whose makers differ from their servers (for example, an employer may make a policy that applies to his employees: 'all employees must write weekly progress reports').

6 References

[Bedford-Roberts 91] James Bedford-Roberts, *Policy Management in Pythagoras*. Technical report, Version 1.0, January 1991.

[Charniak 85] E. Charniak and D. McDermott, *Introduction to Artificial Intelligence.* Chapter 9. Addison-Wesley, 1985

[Moffett 88] Moffett J.D. and Sloman M.S. *The source of Authority for Commercial Access Control*. IEEE Computer, vol 21 no 2, pp59-69. Feb 1988.

[Moffett 90] Moffett J.D. Private communication. December 1990.

[Randell 91] Jim Randell. *Python, a language for the representation of policy*. Technical report, Version 0.1, January 1991. Hewlett-Packard Laboratories, Bristol.

[Sadler 91] Martin Sadler. *Policy Representation*. January 1991. Technical report. Hewlett-Packard Laboratories, Bristol.

7 Glossary

- **ability** an *effect* accomplishable by an *individual.*
- **action** an *effect* accomplished by an *individual.*
- **atomic effect** an *effect* considered to have no *sub-effects.*
- **borrower** an *individual* that gains permission to access another *individual's possession.*
- **client** an *individual* with a *right.*
- **commitment** the intention of an *individual* to act with integrity towards another.
- **cooperation** a set of *delegations* balanced so as to be beneficial to all the *individuals* concerned.
- **cooperative arrangement** a collection of *rights* and *responsibilities* associated with a *goal.*
- **delegation** an arrangement whereby one *individual* pursues a *goal* of another *individual.*
- **effect** a *perception* of a state change.
- **entity** an atomic unit, something that may be viewed as a whole.
- **enterprise** a group of *individuals cooperating* to pursue a set of common *goals.*
- **event** an *effect* that has occurred.
- **goal** an *effect* desired by an *individual.*
- **hat** a collection of *rights* and *responsibilities* associated with a single *individual.*
- **head** anything that assume (wear) the rights and responsibilities associated with a *hat*, it may be a person or a group of people with a constitution to bind them.
- **individual** an *entity* capable of having perceptions.
- **insatiable goal** a *goal* with no associated condition of accomplishment.
- **owner** the *individual* in an *ownership* relation.
- **ownership** a perceived relationship between an *individual* and an *entity*, whereby the *individual* is believed to have ultimate control of the *entity.*
- **policy** a *right* or *responsibility* declared so as to prompt *action* conformant with an intention. (See section 5.4.)
- **procedure** an organized arrangement of *abilities* associated with an *individual.*
- **property** the *entity* in an *ownership* relation.
- **responsibility** marks a *goal* held by an *individual* that has been delegated from another.

- **right** marks a *goal* held by an *individual* that has been delegated to another.
- **role** a set of *rights* and *responsibilities* associated with a single *individual*.
- **root effect** an *effect* considered to have no *super-effects*.
- **server** an *individual* with a *responsibility*.
- **state** the condition in which an *entity* is.
- **strategy** an organized arrangement of *abilities* associated with a *goal*, such that the *goal* can be accomplished by the orderly accomplishment of the *abilities*.
- **sub-effect** an *effect* resulting from the decomposition of another *effect*.
- **super-effect** an *effect* resulting from the composition of other *effects*.
- **system** an *entity* conceived by an observer to delimit the extent of what he wishes to observe.
- **trust** is the belief of an *individual* in the integrity of another.