Collaboration as an Activity

Coordinating with Pseudo-Collective Objects

David Zager, Ph.D.
Avesta Technologies, Inc.
2 Rector Street
New York NY 10006
212-209-1505 (v) • 212-285-1551 (f)
dza@avesta.com

Abstract

A *coalition* is a collaborative pattern in which people must work together to accomplish a task, but where organizational constraints stand in the way of their making use of the coordination techniques that typically enable collaboration. Trinity is a software system that uses a virtual world to provide coalitions with synthetic coordination capabilities functionally equivalent to those that occur naturally in an organized collaboration. The virtual world functions as a *pseudo-collective object* since it plays the coordination role of a collective object. The design philosophy underlying Trinity is heavily informed by some of the fundamental concepts of Activity Theory. To illustrate these ideas, we will examine below the case of coordinating operational information for the coalition of people who maintain the stream of data coursing through a securities brokerage, describing along the way the relevant Activity Theory background and the architecture of the virtual world.

1 Design challenge

The loosely organized collaborative pattern I call a *coalition* offers few coordination services to its constituent collaborators, and the collaboration suffers accordingly. A common feature of the coalition is its inefficiency – participants duplicate each other’s efforts, and many problems often fail to resolve either quickly or to anyone’s satisfaction.

Managing an environment where coalitions abound demands coordinating the coalition through synthetic means. The approach explored here is to use a mediation artifact in the form of a *virtual world* whose state is tightly coupled to the real world. This virtual world provides constituents of the coalition with an account of their world: that is, it situates phenomena of the external world into the context of “the world as constituents see it,” and allows constituents to associate their perspectives with those of others. Extending the terms of Activity Theory, I call this artifact a *pseudo-collective object*.

The design of a prototype system to support coalitions through a virtual world is based on the author’s experiences in managing the systems underlying brokerage operations and an Activity Theory analysis of the domain. The application that embodies the virtual world is a commercially available product known as Trinity. Trinity has many uses in systems management, most of which are purely practical in ways that are not particularly geared to enhancing collaboration. The system has been running in full production mode for only a short time, and field results are trickling in. So far, they have revealed more about the product's lower-level tactical operations functions than about higher-level management goals such as collaboration.

1.1 Collaboration without collective objects

Coalitions are temporary collaborative groups where shared concerns are the filaments connecting constituent individuals and teams. Constituents are part-time members of the coalition, making the coalition loosely bound. At any moment, the coalition's membership is fluid and diffuse, and communications among constituents may be non-existent. Coalitions lack the semiotic self-regulation (Raeithel, 1998) of typical collaborations, where team members coordinate their activities through talking to one another as well as through interacting with their tools. The coalition need not have shared a collective object to set it in motion, and may not ever have created one during the course of its existence.

Research in collaboration, whether from situated activity, distributed cognition or Activity Theory perspectives, tends to focus on the coordination techniques of people working together, thus assuming environments where people share information. For instance, situated activity research (*e.g.*, Suchman 1998; Goodwin & Goodwin 1998) sheds light on how working practices and interactions among participants are critical to coordinating the work environment. Distributed cognition (*e.g.*, Rogers & Ellis 1994; Hutchins & Klausen 1998) shifts the focus to functional systems and the coordination techniques of people working within them, where a functional system is a collection of individuals and artifacts and their relations to each other in a particular work practice. Activity Theory (*e.g.*, Engeström 1998; Engeström, Engeström & Vähäaho 1999, Nardi 1999) emphasizes how the dynamic structure of participants' interactions generate the functional system, rather than the functional system itself. Despite the differences between these approaches, the typical cases under study are those in which conversation or other interaction among co-workers is a normal and essential ingredient to their successful collaboration because it provides coordination of their joint efforts. Co-workers naturally create an amalgam of perspectives of the workplace, a world of shared languages and related tasks.

When such interaction does not exist, as is the case of the coalition, co-workers find it difficult to coordinate with one another. Our hypothesis is that it is possible to manage a coalition by creating a pseudo-collective object. This object can fulfill many of the information-sharing functions that collaborative interactions do, and thus provide a coalition with coordination services that are functionally similar to what the collective object provides a typical collaboration. The structure of the pseudo-collective object is a complex set of interrelationships among objects. Its function is to situate individual events into a larger surrounding context – what in Saussurian semiotics would comprise a system (Saussure 1965). In so doing, it creates an artificial amalgam of conceptual perspectives on the workplace.

1.2 Intensional and extensional perspectives

It is useful to distinguish the extensional versus the intensional perspectives of a model of some world. A *model* of a system is the combination of some organizing principles, the structure of the system, and the elements or objects making up the system. *Extension* focuses on the relationship between some object in the model and a corresponding thing in the external world. *Intension* focuses on the meaning of an object in the model owing to the system of relationships in which it participates with other objects. To be meaningful – to relate to the information consumer's worldview – an object in the model must ultimately have both an extension and an intension. In general, it is necessary to understand which thing is the subject of conversation, as well as what significance the thing receives in the context of that conversation.

The notions of system and the interrelationships among components are critical. Let us take a concrete example of a game of Bridge in which I am holding the Two of Clubs. In this example, the rules of the game are the operating principles; the cards themselves are the objects or elements making up the system; and the logic underlying the system of cards is the structure of the system. I can describe that Two of Clubs physically or logically, and be referring to an individual card in the world (the one that is in my hand). This simple referential statement, however, fails to note the significance of the fact that in this particular game of Bridge, the Two of Clubs is the highest outstanding trump; that I have a void in Diamonds; and that the player to my right led a Diamond in this trick. It is possible to interpret the significance of the card only when taking into account an understanding of the rules of the game, the initial distribution of cards, and the state of this particular game relative to those rules and that distribution. The purely referential description fails to address the significance of the card relative to the system of the game.

Collaborations rely on collective objects to impart significance on the events and conditions their constituents encounter. Coalitions lack collective objects, and hence their constituents may not be able to see the significance of events, only their referential descriptions. Consider, for example, a coalition of people associated with maintaining the operational status of networked computing systems. A member of that coalition may be a systems administrator who becomes aware that "that beat-up old printer on the 15th floor crashed again" (extensional perspective), yet does not see that the printer failure leads to "no paychecks today" (intensional perspective). The role of the pseudo-collective object is to furnish members of coalitions with the significance of the things in their world – to arm that system administrator with the information about why the printer failure matters and allow him or her to set priorities accordingly.

1.3 Seeing the issues in terms of Activity Theory

A recent article of Davydov's (1998) articulates the major challenges that Activity Theory faces in his estimation. Prominent among them is the question of whether the collective object and subject exist – and if they do exist, what they are. I will argue that collective objects are sign systems that function as mediating artifacts (Engeström 1998) with which groups of collaborators coordinate themselves. In other words, the collective object is a shared, socially constructed system of interpretation of the external world.

Collective objects coordinate people working in disparate functions by providing a context beyond the individual. Work tasks can be seen extensionally to refer to some action, some aspect of a job that has independent existence in the external world. For instance, suppose my job is to look for reported faults in router interfaces, and it comprises watching network management consoles for alarms for extended periods of the day. In a broader organizational context, the reason for my actions is to ensure the quality of a data feed that others rely on for performing their actions.

Suppose I notice two identical faults in physically identical router interfaces that occur within seconds of one another. To understand the situation from an extensional perspective is simply to recognize the similarity of the two faults. Understanding the situation from an intensional perspective demands I know that a fault in *this* router interface stands in the way of trading (on which the organization places a high value), while the other failure hampers delivery of supplies to the building (on which the organization places relatively low value). In other words, in the intensional perspective, local phenomena are interpreted in terms of the rules of the game: the situating context helps assign a value to otherwise indistinguishable referents in this interpretation. The collective object of the context is a contract that the participants have jointly created. This contract has given the participants a common vocabulary and interpretation of the world for the duration and extent of a group activity.[1]

When participants lose sight of their roles in the larger context of the game, other problems arise. Devoid of a broader context, priority becomes a function of inherent interest or severity of the problem, which is often immaterial: a boring and simple problem may prevent many other people from accomplishing their tasks, while a theoretically fascinating problem may not make one iota of practical difference. Lacking the means of distinguishing issues to deal with immediately from those that can wait until later, participants without context may end up with skewed priorities and take inappropriate actions. Managing coalitions demands addressing these issues.

Accounts of collective activity tend to focus on the role of the individual, despite account writers' intentions of describing what goes on in a group. Consider, for instance, how in *Problems of the Development of the Mind*, Leont'ev (1981) wonders about how it is "possible for action to arise, *i.e.*, for there to be a division between the object of activity and its motive. It obviously only becomes possible in a joint, collective process of acting on nature." (p. 210) He arrives at this point by examining "the fundamental structure of the individual's activity in the conditions of a collective labour process. (...) When a member of a group performs his labour activity he also does it to satisfy one of his needs." He finds that the "product of the process as a whole, which meets the need of the group, also leads to satisfaction of the needs of the separate individual as well, although he himself may not perform the final operations (...) which directly lead to possession of the object of the given need."

The example he elaborates considers the distribution of labor in a hunt. An individual hunter needs food. His action in the activity of hunting aimed directly at beating the bush to frighten a herd of animals and send them to a group of hunters waiting in ambush. Leont'ev observes that the activity of this individual member of the hunt ends with this action, although the result of the action does not in itself, and may not, lead to satisfaction of the beater's need for food. There was a lack of coincidence between what the processes of his activity were directed to (*i.e.*, the motive of his activity) and what stimulated them. He goes on to define *action* as the case when the object and motive of a process do not coincide with one another. In the current example, the beater's *activity* is the hunt, and his *action* the frightening of game.

This is the classic separation of action from activity, with the further distinction of goals, needs and objects (see also Davydov, Zinchenko & Talyzina 1982). The hunt is an activity; its *object* is dead animals; the hunter's *need* is food, which he hopes to satisfy by participating in the hunt. The hunter's participation itself is goal-oriented – his goal is to beat the bush to roust the game from it, and to ward that game into the ambush. Leont'ev observes that what unites the direct results of this activity with its final outcome is nothing other than the given individual's relation with the other members of the group. For Leont'ev, the objective basis of the specific structure of the individual's activity is the activity of other people, leading him to conclude that historically the connection between the motive and the object of an action reflects objective social connections and relations rather than natural ones.

Frightening the game satisfied the goal of the hunter's action in the activity of hunting, though not his need. According to the division of labor, the hunter now hands over responsibility to the hunters in ambush. The hunter will not realize his need to get food at least until the group's activity of hunting has realized its object: an adequate number of dead animals.

Leont’ev’s emphasis is on the individual’s object here, despite the fact that he is talking about a collaborative effort. He provides neither vocabulary nor framework for exploring the collective object of the group of hunters. The reason for this is that he has concentrated on the extension of the object rather than its intension. What matters about a collective object, however, is its intension more than its extension – the way that different individuals’ needs interact to form a system rather than the way they refer to things in the world. The organization of a collective activity like the hunt is a sign system. It is a mediator between “raw” data of the world and the hunters’ conceptualization of the world (Vygotsky 1987, 1978). The intension of the collective object is the set of participants and their dependencies and interactions, *i.e.*, the mediating artifacts that shapes the collaboration. The hunter, that is, is not simply beating the bush: his beating is contextualized by the larger situation of hunting – the set of structured relationships among participants, and the relationships between the participants and the environment in which and with which they interact. Participants define and interpret their own and others’ activities, as well as the world outside, in terms of the system of interactions, *i.e.*, the hunt.

On the other hand, the extension of the collective object is little more than an abstraction of what participants want from the process. The single outcome of the hunt matches the object, not the needs of the hunt. The object, dead animals to satisfy all needs, is true regardless of individual need. Conversely, divergence of needs among participants is irrelevant to the collaboration of the hunt. It makes no difference that your need in the hunt is sinew for some cool tent-post straps, while mine is to satisfy a craving for gazelle salami. The collective activity subsumes the individual. The potential for meeting the individual's need comes from the object of the collective activity, not the object of the individual's activity. The collective of hunters shares a single object in its activity. That is, assuming that "collective object" is the same as the "object of the collective activity," the equivalence is not

collective object = collection of individual needs,

but rather,

collective object = single outcome.

All participants need game for their own purposes. They can get it only by banding together. Therefore, they band together, divvy up the spoils, and all participants satisfy their own needs to the extent it is possible.

When groups of people working together have no means of creating such a collective object, their coordination suffers. If we apply these notions to the complexities of the current high technology work environment, we propose that a software solution that offers a synthetic picture of the context, and captures the intension of collaboration in terms of the service interrelationships of the computing components, will be able to mitigate some of the effects of lack of coordination.

2 Architecture and implementation of the approach

The particular coalition whose interests we are trying to satisfy (and that we will describe in §3 below) is the set of people with operational responsibilities for the networked computing environment of a large global securities firm whose systems offer complex and heterogeneous services. In this context, the purpose of the virtual world is to provide each constituent or group of constituents in the coalition with the information they need about the state of resources in the world as they see it. That is, it presents information according to each audience's basic level of categorization (Rosch et al, 1976), and limits the immediately available information to the resources that any individual decides are within the range of interest. The virtual world structurally situates the individual's microworld into the larger context. Its calculations show the viewer not only the structural relationships among things he or she is responsible for, but also how they relate to other components of the system, both upstream and downstream. As a static description, this may not be interesting to the information consumer. Dynamically, however, this set of interrelationships among networked computing components is descriptive of the implications of faults through the network, both *horizontally* – between components at the same layer of processing – and *vertically* – or between layers of the network stack running from the physical network through the presentation portion of an application. The set of interrelationships that hold among resources as part of the play of the game (*i.e.*, the intensional perspective of the network) is predictive of the way that the impacts of faults move through the network. These same interrelationships, therefore, are able to convey interruptions of service, and are thus instrumental in managing problems that arise in operating the various facets of networked computing systems.

2.1 Working definition of a virtual world

The software embodiment of the pseudo-collective object is a mediation tool based on a virtual world. This virtual world shares many of the defining characteristics of Multiple User Domains ("MUDs," Curtis 1992, Curtis & Nichols 1994), namely, that it is an *interactive*, *multi-user*, *text-based* virtual reality. The interactions among the users and the stable (or more-slowly-changing) organization of the surrounding environment create the "worldness" of the virtual world. The mediation tool differs from most other virtual worlds in a critical way, however. The typical virtual world is a simulated environment in which participants interact virtually, and independently of the external world. A simulation, that is, starts with seed values from the outside, and progresses of its own accord, executing its own internal logic. After some elapsed time, the observer can (if desired) compare the state of the simulation to possible states of the external world, and determine whether the simulation is an accurate model of the world. In the current case, however, the design of the virtual world (Zager 1999) is *not* a simulation. It is autopoietic (Maturana & Varela 1973) in nature, meaning that the virtual world stays structurally coupled to the world external to it, and hence the virtual world reflects changes to the state of the external world with commensurate changes to its own state.

In addition to seeing an inventory of things in the world that are of interest to them, observers also receive a causal account of state changes among those things. At any moment, both the state of the virtual world itself and the states of the interconnected components making it up are reflective of the state of the external world. Because the relationships within the virtual world imply dependencies among resources in the external world, the viewer has the illusion of effective causality in the virtual world. That is, the virtual world's calculations relate underlying change in one object (in this context, typically some fault) to visible effects in others.

The virtual world contains a big picture of information. That big picture is so abstract, so replete with heterogeneous information, and so detached from the situation of any individual functional system that it is not useful in and of itself. The "big picture" description exists solely as a logical consequence of the virtual world's containing heterogeneous information. It makes no claims about cognitive needs or abilities, and does not imply that any human observer sees or is aware of the whole system. Only individualized projections can convey information meaningfully. Each audience needs see only those parts of the virtual world that are salient to its concerns, but is capable of seeing how its areas of interest interrelate with other phenomena that are, under natural conditions, beyond its scope of vision. The Goodwins' observation (Goodwin & Goodwin 1998) that different individuals' "seeing" airplanes varied according to the work they are engaged in holds just as true for members of a coalition as it did for a coordinated work environment. The purpose of the virtual world is to allow each member of the coalition to "see" the world in a different way by enabling each participant to construct individually tailored projections corresponding to his or her basic level of categorization of the world, and to make that view the default perspective for that individual. The ability to project any participant's perspective of the world demands that the tool contain a superset of all perspectives that are of interest to the participants of the coalition – and hence the illusion of a big picture. It is important to note that despite the visual implication of the terms of the ensuing discussion, such as *projection*, *observer*, or *narrative histories*, the emphasis of the discussion concerns the structuring of data, and the conceptual organization of data that some graphical display might then present, rather than the graphical display of information. In the end, innovative presentation techniques will be critical for effective interaction with the product; they are simply not our focus here.

2.2 Brief architectural description of the virtual world

The major components of the system, illustrated in Figure 1, form a three-tier architecture, comprising

- data collection and analysis,
- the virtual world, which functions as an active database of highly interrelated objects, and
- user stations that allow the information consumer to interact with customized projections of the information in the virtual world

<INSERT FIGURE 1 HERE>

2.2.1 Data collectors

The data collectors have two primary functions: i) discovery of the things in the world around them (network devices, computers, applications, people) and the interconnections among those things; and ii) discovery of changes in the condition of those things or their interconnections. They can discover objects as small as an interface on a router, or as large as a business group (which it may discover by reading a human resources file, for instance), and all points in between. In a large enough organization, this means that data collectors discover many thousands of objects of a wide variety of types, as well as the relationships among those objects. Discovery of a new object amounts to an insert operation for an active database; a report of a change in condition is the equivalent of an update operation.

With respect to the overall system, we may say that the data collection layer's responsibilities are extensional: it needs to "know" referential descriptions only, to recognize explicitly engineered components according to their tagged identifiers (such as variables in a Management Information Base or MIB), with no regard to the significance of the things it collects data about. Consider the sample network fragment in Figure 2. When the data collectors identify a networked desktop computer, they insert it into the virtual world as a physical node that embodies an operating system and other programs, including a TCP/IP networking stack, a Network Interface Card ("NIC"), which has a Media Access Control ("MAC") address. The NIC connects through a cable to a hub, and ultimately to a corresponding port on a router, which has its own MAC address. Each object – NIC, TCP/IP stack, MAC address, physical node, operating system, applications, people using the applications – is an entity in the active database. The database does not contain just an inventory of types arranged in a hierarchy, however. It also contains the structural and functional relationships among the objects, which are known as *services* in the language of the virtual world. In that way, the physical node provides an *embodiment* service to the operating system, the NIC provides *networking* service to the computer, and so forth.

<INSERT FIGURE 2 HERE>

2.2.2 Virtual world (active database)

In contrast, the virtual world's responsibilities in the overall system are intensional: its function is to "know" about the significance of objects, focusing on the interrelationships among objects rather than the things the objects refer to. A consequence of this change in emphasis is that the process of preparing an object for the virtual world strips it of most of the attributes that the data collection layer had needed to identify things as themselves. What remains is a set of objects that can be distinguished from one another by the set of relationships each participates in. The semantic abstraction of the virtual world emphasizes the meaning of objects owing to their positions within a system of relations rather than to any intrinsic meaning to the objects themselves. Seen from another angle, all knowledge in the virtual world is local. Each of the thousands of individual objects "knows" only the list of objects it provides service to, and the list of objects it consumes service from, where, as noted above, the term *service* describes a relationship between objects. Services are recursive (*i.e.*, services can be built on other services), and become increasingly abstract and complex as we move from the fundamental components towards the needs of the information consumers. Thus, a securities trader's *trading* service may comprise: voice line(s), current tick data, historical data, news feeds, modeling tools, and a trading application. A combination of the desktop workstation, a number of servers, data feeds, network connections, network devices, databases, and miscellaneous other components collaborate to deliver that service. When one of the components is not working properly, the service as a whole is degraded. The virtual world seeks to represent the set of relationships that make up that service, as well as the effects on that service that faults in any of its component objects may bring about.

Although an observer of the virtual world might think of it as containing a unified picture of the external world as a whole, in fact no component of the system holds such a picture. The discovery of the external world through its extensions has constrained and defined the possible sets of interrelations among the elements, freeing the virtual world itself to focus entirely on the intensional perspective.

The virtual world's function in terms of processing is to calculate the changes to the overall system that occur given information about the change in state of any element of that system. It accepts information of changes to the external world – either the insertion of new elements, or changes of state to known elements. Each object that changes owing to external information informs its list of service partners about its changed state. Each object that receives such an internal message determines for itself whether or not its partner's change of state affects it in any way. The way that objects indicates that a service provider's state change matters to it, is to change its own state accordingly, and to create an association between the set of state-changing events it owns (called an *alarm*) and the set that its service provider owns. The net effect of this set of information exchanges within the virtual world is a set of associated alarms *corresponding* to the chain of objects with associated state changes. This set of sets of events (or *episode*) thus draws together apparently independent phenomena in a way that is constrained by the structure of service dependencies in the model, but is defined by the actuality of a particular set of interrelated changes. An episode represents the life history of an event in the external world, consolidating the reports of changes that have emanated from possibly a multitude of different data collection systems. In the example of the Bridge game above, I argued it is possible to interpret the significance of a card only when taking into account an understanding of the rules of the game, the initial distribution of cards, and the state of this particular game relative to those rules. In the same way, it is possible to understand the significance of an object's change of state only when knowing:

the rules of the game = how the network fits together

the initial distribution of cards = the functions people are trying to accomplish with this particular network

the state of this particular game relative to those rules = how the current availability of components relates to the ability to execute those functions

For a more concrete example of an intensional perspective in action, consider Figure 2 above once again. The path from an Accounting workstation to the printer can go through either of two routers, A and B. If A malfunctions, it does not matter because print requests can still flow through B; similarly if B malfunctions but A remains active. The significance of losing router A to an accountant's ability to print derives from interpreting the state of the larger context at the moment of asking the question because it depends on whether or not B is active. Through its construction of episodes, the virtual world constantly evaluates the significance of events. It notifies observers of the system when it calculates that that significance coincides with the observers' sets of interests.

2.2.3 Presentation layer

Given a large enough organization that provides a plethora of different high-level consumable services to a diverse community of workers, understanding the set of objects within the virtual world and the set of interrelationships among them can exceed any individual's grasp. In reality, an individual information consumer needs typically to understand only a tiny proportion of the whole – for instance, what goes into providing trading services for this particular set of traders – that are projections of the overall model. Consequently, the projections carve up the high-level picture, and present it at the proper basic level of categorization for any particular audience. Since these are all simply different views on a single model, any projection can transform to any other through navigation – meaning that the virtual world provides the *lingua franca* that allows for intertranslatability among worldviews.

The mechanisms for creating such projections are called filters. Filters are similar to the notion of views found in relational databases. A filter is a grouping of resources governed by a descriptive language of inclusion criteria, in similar spirit to Malone et al. (1995, *cf.* also Malone et al. 1986). The simplest filter is a specific list of objects within the virtual world. Filters that are more complex comprise expressions that describe characteristics of the objects they contain – for instance, all commodities pricing feeds, or all databases located on the 34th floor. The virtual world constantly reevaluates its expressive filters, meaning that as new objects are inserted into the virtual world, or existing objects change their characteristics, they will become grouped by whatever preexisting filters match their characteristics.

Filters group resources according to both horizontal and vertical criteria. A horizontal projection might result in, say, application resources; a vertical projection might run the gamut from the applications running for the trading desk, down to the router interfaces allowing data to reach that desk. In general, such filtering brings to the fore those resources that correspond to the information consumer's basic level of categorization of the world. Because the model underlying these resources is a highly interconnected graph, the presence of each resource implies the presence of the resources with which it has connections. Thus, the subgraph of modeled resources underlying the view of a router interface may be the same as that beneath the view of a trader's application, varying in the particular resources promoted to salience in the consumer's interface.

The end user interacts with the virtual world through filters, not directly, meaning that the filters always provide a mediated view of the virtual world. Interaction is primarily of two sorts: unsolicited reports of changes, and exploration. By constructing or selecting filters, the information consumer is requesting that the virtual world inform him or her of any change to the filter – new resources added to it, resources removed from it, or changes to the state of resources it contains. The virtual world is responsible for alerting the information consumer of such change. The user station can present the fact of change through textual interface (*e.g.*, the text of an event that caused the virtual world to change the state of some resource) or through graphical interface (*e.g.*, changing the color of a graphical depiction of a service relationship to indicate changed state). If a service relationship has changed, the information consumer may want to understand why. The virtual world user interface gives the consumer the ability to decompose services into the components that allow that service to be offered – in other words to traverse the underlying dependency graph of the virtual world.

Different visualization techniques are appropriate for different users. Some prefer to see the audit trail of events that have led the virtual world to change the state of one of its objects; others prefer to see a visual depiction of the objects themselves and their interrelationships. Since the virtual world contains all this information, these techniques are only superficially different presentation mechanisms of the same underlying data projection.

3 Example of a coalition

An example of a functioning coalition is the set of organizational structures that contribute to assembling and maintaining the ensemble of networked computing components and services that support brokerage operations in a modern global financial institution.

Operational challenges confronting that environment include:

- increasing complexity of the software and hardware systems themselves
- increasing fragmentation of knowledge about system components
- increasing rate of change of software and hardware components
- increasing emphasis – and reliance – on information technology services in conducting normal everyday business process
- increasing decentralization of control of information technology services
- vanishing supply of staff who were skilled enough to manage either the complexity or the rate of change, and were experienced in *this particular environment*
- lack of tools to help meet these challenges

A common thread of buying and selling of securities runs through much of brokerage operations. These securities are specified through a tuple-space that includes: security symbol and identifier, market it is traded on, price (and currency), time of a trade, bid and ask prices at the time of the trade. A steady stream of information known collectively as “market data” enters the securities firm from a combination of third-party providers and direct feeds from exchanges around the world. These data provide information to brokers and traders at the time of a trade; to analysts researching the performance of a company; to people constructing complex deals or contracts; to people who need to clear and settle trades after they occur; to people who need to provide custody of the underlying securities; to accountants who need to determine the price of a deal, including those that span multiple currencies; to people who prepare customer statements; to people who provide real-time positions monitors for their customers; and so forth. To make market data available to these information consumers, some of the subject matter experts ("SME"s) who must be involved are:

- People with expertise in network engineering and network operations. The stream of incoming data at the open and close of a market can be substantial. It must be distributed widely through a firm, and in a timely manner.
- People with expertise in systems administration are concerned with preparing server-class machines to cache incoming data for multiple users, and for ensuring that the data consumers' desktops are well-tuned and responsive.
- People who are applications specialists normalize incoming data, since different exchanges and different third-party feeds represent the panoply of possible data points in non-standard ways. Additionally, not all securities are analyzed along the same dimensions, yet further use demands the ability to compare them, which is possible only after normalization.
- People who are specialists responsible for inspecting the contents of incoming data, for noting any new symbols that appear in the stream (new options created,

new companies, mergers and acquisitions, etc.) and adding those to a product database, for recording “corporate actions” such as dividend payments, and so forth.

- People who are concerned with storing market tick data (changes in bid- and ask-price for securities) in time series databases, aggregated and normalized for fast recall.
- People who synchronize tick databases with incoming “pricing feeds” that provide the information on actual prices of trades as they occurred on specific markets (since the price of, say, IBM on the New York market may not be exactly the same as the price on the Boston market). The databases must also synchronize with and integrate the fluctuating foreign exchange pricing feed, to calculate the US dollar value of a trade made in Italian lire and settled in Japanese yen on a London exchange.
- These people, in turn, rely also on systems administrators for maintaining their servers and watching over the vast farms of disk storage needed to house the ever-growing mountain of data.
- Data quality experts who need to get involved with checking the quality of information coming in on all feeds, using whatever heuristics available to them to look for possible anomalies and correct them. Pricing feeds and market data feeds are prone to data exceptions, owing to human error, noise on communications lines, and other factors.

- Middleware specialists who allow the data of incoming feeds to distribute across the various interest groups of an organization, whether synchronizing redundant feeds entering the organization simultaneously in New York and London, or spreading the data of a single Paris feed to traders in New York, Frankfurt and Singapore.
- Specialized application groups who incorporate the data of the feeds into their applications–from trading and modeling applications through real-time options calculators, portfolio calculators and position monitors, the applications that associate prices and commissions with trades, analytic applications that construct trends in securities performance, and so forth.

This list is not exhaustive, simply representative of the kinds of expertise that are necessary for delivering market data within a financial services firm. It leaves out real-time news, historical news, TV and radio broadcasts available online, and many other aspects of data delivery.

Information services organizations have traditionally separated operations from engineering (programming, network configuration, etc.), and network management from both systems and application management. Continuing trends in outsourcing of commoditized functions accentuate the existing separation, since paycheck-loyalty now is distributed among many different employers. Application hosting for back-office functions, such as accounting and human resources, accelerates the organizational entropy.

Yet, when a breakdown occurs at a deeply embedded layer of this environment, such as an interface failure on a network switch, every one of the SME groups, and their respective information consumer communities, feels the effects in terms of the interruption to the consumable services they are responsible for. Whether or not the groups are aware of one another's existence or function, the complex web of networked services binds the people who have operational responsibilities for delivering those services into a set of dependent relations. Different breakdowns affect information consumers differently. For instance, loss of a switch interface may have no effect on business units with fully redundant network circuitry, and therefore be an issue only to network operations personnel. If an earlier failure had interrupted the backup line, however, this same loss of a switch interface could pose a major disruption of service. The interface is the same in both cases, the failure itself is the same, but the interpretation of the failure in context is different owing to the interplay of many factors, parallel to the example of the Bridge game above. The effect on the community of SME groups is that the set of personnel who must deal with any problem changes according to the particular context of the issue problem, not its underlying referent.

3.1 Coordination of a coalition

Together, the SME groups affected by a problem form a coalition. The organization of the coalition is bottom-up, comprising independent participants acting on their own, with little or no reference to the other participants. The active environment with which they all interact, and which coordinates the collaboration, comprises the distributed computing system itself, particularly its faults.

Consider again the example of networked computing operations in a large organization, looking at the simple network fragment shown in Figure 2. Suppose Router B is powered off and a card containing multiple ports fails in Router A. As a result, connectivity between LAN A and LAN B is lost. To the end users, this fact is visible in terms of a number of symptoms:

- Several users' workstations receive no response from the mail server.
- It is impossible for the traders to reach the Market data server, leading to a halt in the live feeds.
- Router A broadcasts failure messages via SNMP (Simple Network Management Protocol – a standard mechanism for networked devices to report self-diagnostics).
- Local printer queues become backlogged.

In other words, Internet connectivity to half the enterprise is lost, and critical business services are interrupted. Many competing viewpoints categorize the network and diagnose the problems that arise within it differently, as Table 1 illustrates :

This group …	**… sees a collection of …**	**… categorized by …**
Network managers	nodes	subnetwork (addressing scheme)
Network administrators	nodes comprising replaceable components	physical interconnectedness (network topology)
System administrators	computing nodes and peripherals	users and office layout
Application specialists	application instances	code level and component distribution

This group …	… sees a collection of …	… categorized by …
Data domain experts	services and client communities	service and data usage
Database admins	databases, data, servers and storage devices	users, data sources, demand patterns
Data quality experts	time series	feed sources
Middleware specialists	interconnected messaging applications	geographic location of user and data source
Users	apps, workstations and peripherals	consumable services accessible to a person
Business unit managers	people and business services	financial ownership
Help desk staff	users and services	users or services

Table 1: Different perceptions of the network

For instance, if I am a network operator, I see the business unit as a set of switches, routers, repeaters, external lines, and so forth. On the other hand, if I am a business manager, I see that same business unit as a set of individuals with differing responsibilities and changing positions. If the network components fail, the business services fail. By default, the network operator wants to see what network component failed; the business manager wants to see which staff members have lost their access to information. Either the operator or the business manager may want to explore beyond the initial view: the business manager might want to look at the underlying system problem, while the network operator might want look at what business unit has suffered the brunt of the problem. Under normal working conditions, they probably do not stray that far from their basic level of categorization.

We can depict a breakdown in a switch interface in terms of different perspectives on a single underlying system fault, as in Figure 3. The DBA will try to fix the problems that New York customers have in reaching a San Francisco database. A data quality expert in London finds discrepancies in a time series for a tick database, and attempts to repair the data. Application support in New York finds overloading on a number of servers that offer remote computational services via CORBA when they are expected to be redundant overflow machines and attempts to sort out load balancing manually to improve system responsiveness. None of these SMEs conducting troubleshooting on a purely local level is at all aware of the trap (unsolicited report of an error coming typically from a device) that the network operator has seen coming from the switch card, just as the network operator has no inkling of who uses the circuits that pass through that card.

<INSERT FIGURE 3 HERE>

The participants in the coalition *are* collaborating, albeit inefficiently. Each participant contributes uniquely to producing a conjoint product – restored homeostasis of computational services. They coordinate their individual activities through their reactions to the evolving state of faults seen locally. Whether or not the SMEs are aware of it, there are structural and functional interdependencies binding them together. When the network operator reseats the card and the switch starts functioning again, the conditions are fit for resolving the other problems. Some problems will resolve automatically, others will take continued specialized intervention.

What differentiates the functioning of the coalition from collaboration as an activity, however, is the lack of a collective object serving as the mediating artifact that enables the participants to interact. As SMEs react to the changes to the aspects of the networked computing systems that they perceive, the set of direct actions they take to restore service gives rise to any coordination of the coalition. The changes in the surrounding computing systems, that is, function for the SME as the semiotic indicators of how to react. They convey a completely localized message of what has failed in the systems the SME is responsible for, and so have transformed into the medium through which the SMEs learn of issues. Since, however, different characteristics of the systems appear differently depending on the SME's set of responsibilities, perceived changes make for an extremely noisy and inefficient communications medium. They do not coordinate one SME's local conditions with another's, and therefore lead to needless duplications of effort, and to addressing of symptoms rather than of operationally prior underlying causes.

4 The pseudo-collective object

Collaborators who coordinate with one another through conversation accomplish that coordination by providing each other the necessary information about how the collaborative work is progressing, translating their *account* into common vocabulary. Suppose the network operator had informed each SME in turn, for instance, "Card number 7 on switch AQ9036W has failed, so you should expect to see slowdowns in the rate of market data delivery / missing data in historical databases / connection failures in remote options calculators / …" and therefore took some of the guesswork out of local troubleshooting. Each SME who had received such an account of the outstanding problems could have focused on alerting the end users, and on coordinating their efforts of restoring service with the state of either the switch card or some workaround.

Not only people, but systems, too, can provide accounts, as long as their designers have built them in. Dourish (1997) observed that when a typical user of a system is trying to get some work done, an account of state allows that user to tie the system's progress to the user's own, where an account is "the story the system tells about itself – its presentation of its own operation and state (and the relationship between the two)." Dourish's example pertains to photocopiers, where the typical audience of an account is the person trying to make copies. This is not a SME interested in the behavior of the machine for its own sake, but someone who must have his or her copies ready in five minutes, hears strange noises coming from the machine, sees no copies emerging, and is anxious about meeting the deadline. "Why won't the machine tell me what's going on? What do I have to do to meet my deadline?" Accounts that derive from software or from conversation share a common ground: they provide a *narrative history of progress toward some goal*, and offer a *process of information transformation* that enables the hearer to relate his or her own condition to the condition of some other system.

Collaboration that is organized from the top down is a mediated activity (Vygotsky 1987) whose mediating artifacts (Engeström 1998) are conversations, organizational structure charts, contract negotiations, etc. In the operational coalitions we have been considering, the mediating artifact for constituents is the receipt of information from the environment itself. Each constituent encounters operational problems that they interpret according to their own worldviews. That is, if a resource or service that falls within a given SME's responsibility starts showing signs of malfunction or degradation, the SME perceives the effects of that change as semiotic indicators of the state of local conditions. For example, the trader uses an options calculator that runs on a remote server. As far as the trader is concerned, he is using a desktop application. The effect of a failing router interface is to block the path from the workstation to the remote server. As far as the trader and the application support person are concerned, the options calculator that is hanging is the overt manifestation of the interface failure. There is no information immediately recoverable from the problem that ties the visible manifestations back to that failure. The application support person goes through a mental checklist: Is it a disk failure? A problem on the desktop? A network failure? A server application failure? This is imprecise, time-consuming, and, often, not a very successful way of diagnosing the problem.

Purely local evidence fails to coordinate the constituents of the coalition because it loses the interrelationships of changes in the larger environment. The SME does not and cannot interact with "the environment," but only a fragment of the overall surroundings of the workplace. In addition, his or her perception of the environment is not of the things themselves, but rather one aspect of a socially constructed theory of the world through which conceptualization occurs (Vygotsky 1987; cf. Whorf 1964 also). Acontextual interpretation of local conditions isolates the SME, and fosters breakdown of coordination among the constituents of the coalition. The network operator who perceives the problem as a malfunctioning interface has no way to tie her problem to the trader's hanging calculator. The support person on the help desk has access to two unrelated trouble tickets: a hanging calculator and a broken interface, and no principled reason to associate what appear to be independent phenomena. The isolated SMEs lack a recoverable translation from a more comprehensive picture of the world that would normally be available through the coordinating interactions of a collaboration. Setting a wider context with one's neighbors, and understanding one's local interpretation in comparison with others' local interpretations is not the same as advocating a single "big picture." The strength of normal collaborations that we see in the literature comes from their ability to make the separate efforts work together better. When the "big picture" plays a role in this context, it is as an abstraction that is available at best to those who engineer the collaboration; it need have no reality to those who are collaborating.

The virtual world described above can play the role of the missing mediating artifact for the coalition. It provides an account, since it supplies a story for "the system" as if it were telling a story about itself. As a container of incidents across the networked system, it provides the audience of information consumers with a narrative history of progress toward the goal of restoration of service. Because it is possible to create projections of the data contained within the virtual world, the virtual world transforms information about faults according to the basic level of categorization of the consumer. It contextualizes the locally salient aspects of an operational fault to a larger work environment, because it provides a means of capturing the peripheral information that would normally come from the coordinating effects of a collaborative environment. The “big picture” is at least as inaccessible from the virtual world as from a natural collective object: all viewers interact with the virtual world through projections, which mediate their access to it. The projection limits the view to what is understandable *given the viewer’s interests*.

We hypothesize, then, that the virtual world and its projections should function like a collective object for the coalition because it provides the information to coordinate individual efforts, a basis of interpreting multiple accounts, and a narrative history of the progress of some problem the coalition faced. Accounts of the state of computing systems should exhibit the following characteristics to allow for the greatest set of projections:

- encompass the full span of a single fault (perhaps cluster of faults), including its temporal, spatial and contextual variations
    - spatio-temporal variations: portray the life cycle of the problem (episodic analysis) across resources and locations
    - contextual variations: associate apparently unrelated phenomena across resources in the computing systems in a principled way, yielding
        - the source of the problem (root cause analysis)
        - the scope of the overall problem (impact analysis)
- sift out only the information that is of importance to that constituent
- present information to the constituents in terms of their basic level of categorization of the world

- situate the significance of a particular change not only in terms of the constituent’s world, but also the world of the information consumer

With neither top-down organization nor viable interaction among the constituents of the coalition, no actual collective object binds the collaboration into an activity. The accounts that a virtual world provides are pseudo-collective objects. Pseudo-collective objects have all the attributes associated with collective objects – they are mediating artifacts that coordinate the interactions among collaborators and situate them in a larger effort – but they are not the objects of object-oriented activity. The participants in the coalition that is coordinated by accounts do not interact or share worldviews any more than those in the coalition that has no pseudo-collective object. If purely localized accounts of changes in the systems make for an extremely noisy and inefficient communications medium, the pseudo-collective object should serve to make that medium less noisy and more efficient. It has emulated a characteristic of organized collaborations, not transformed the organization of the coalition into a more structured collaboration.

5 Discussion

Collaboration is an activity when people come together to produce some single outcome, when each participant contributes uniquely, and when the results are beyond the capabilities of any individual participant. This joint effort is internally coordinated with a mediating artifact that shapes the way constituents participate. The mediating artifact is the collective object, an abstract notion that refers to the particular sign system through which participants define and understand their own activities as well as the activity of others. It is a system of interrelationships that the individual participants internalize as part of the activity of collaborating, and at the same time they create as part of the activity of collaborating (*cf*. Vygotsky 1978). What differentiates the coalition from the organized collaborative pattern is that the organized pattern is a mediated activity (*i.e.*, it has a collective object) while the coalition is not. Participants in organized collaborative patterns coordinate their collaboration by speaking to one another about progress, creating organizational charts, negotiating with one another and so forth. Such conversations are mediating artifacts, falling within the realm of the collective object, since they are a means of implementing the object, as well as part of the process of creating, maintaining and changing the object.

In both coalitions and collaborations with bottom-up organization, we find the structure of the environment and the structure of the organization mutually constrain patterns of coordination with one another, as the active environment triggers structurally determined sympathetic changes in the organization (Maturana & Varela 1973, also Resnick 1994, Parunak 1997). The difference is that coalitions lack mediating artifacts and hence are not mediated activity.

The members of the coalition diverge in their needs, yet all strive for a single outcome: in the example above, that outcome is maintaining a suitably high level of service for the organization that consumes the services of distributed networked computing systems. When the organizational environment and the computing environment are both highly complex and highly intertwined, the chances of success without coordination are slim. Further field experiments will demonstrate whether a pseudo-collective object in the form of a virtual world can effectively play the role of synthetic mediating artifact.

6 References

Curtis, Pavel and David Nichols, 1994, "MUDs Grow Up: Social Virtual Reality in the Real World," Xerox PARC Technical Report

Curtis, Pavel, 1992, "Mudding: Social Phenomena in Text-Based Virtual Realities," Xerox PARC Technical Report CSL-92-4, available from http://ftp.game.org/pub/mud/text/research/DIAC92.txt

Davydov, V.V., V.P. Zinchenko, and N.F. Talyzina, 1982. "The problem of Activity in the Works of A.N.Leont'ev," in Soviet Psychology 21: 31-42.

Davydov, Vassily V., 1998. "The Content and Unsolved Problems of Activity Theory," In Engeström, Miettinen and Punamäki [eds.], *Perspectives on Activity Theory*. New York: Cambridge University Press

Dourish, Paul, 1997. "Accounting for System Behaviour: Representation, Reflection and Resourceful Action," in *Computers and Design in Context*, Kyng and Mathiassen (eds), Cambridge, MA: MIT Press, 1997

Engeström, Yrjö, 1998, "Activity Theory and Individual and Social Transformation". InEngeström, Miettinen and Punamäki [eds.], *Perspectives on Activity Theory*. New York: Cambridge University Press

Engeström, Yrjö, Ritva Engeström, and Tarja Vähäaho, 1999. "When the Center Does Not Hold: the Importance of Knotworking, " In Chaiklin, Hedegaard, and Jensen (eds.), *Activity Theory and social practice: Cultural-historical approaches*. Aarhus, Denmark: Aarhus University Press.

Goodwin, Charles and Marjorie Harness Goodwin, 1998. "Seeing as Situated Activity: Formulating Planes," in Engeström and Middleton [eds.], *Cognition and Communication at Work*. New York: Cambridge University Press.

Hutchins, Edwin and Tove Klausen, 1998. "Distributed Cognition in an Airline Cockpit," in Engeström and Middleton [eds.], *Cognition and Communication at Work*. New York: Cambridge University Press.

Leont'ev, A.N., 1981. *Problems of the Development of the Mind*. Moscow: Progress Publishers.

Maturana, Humberto R.and Francisco J. Varela, 1973. "Autopoiesis: the Organization of the Living," in *Autopoiesis and Cognition: the Realization of the Living*, Dordrecht NL: Reidel 1980

Malone, Thomas W., Kenneth R. Grant, and Franklyn A. Turbak, "The Information Lens: An Intelligent System for Information Sharing in Organizations," in *CHI '86 Conference proceedings on Human factors in computing systems, April 13-17, 1986*, Boston, MA USA

Malone, Thomas W., Kum-Yew Lai, and Christopher Fry, 1995, "Experiments with Oval: A Radically Tailorable Tool for Cooperative Work," in *ACM Transactions on Information Systems*, Vol 13, No. 2, April 1995, Pages 177-205

Nardi, Bonnie, 1999. "Intensional Networks." to appear in CSCW ...

Parunak, H. Van Dyke, 1997. "'Go to the Ant': Engineering Principles from Natural Multi-Agent Systems," in Annals of Operations Research 75 (1997), special issue on Artificial Intelligence and Management Science: 69-101

Raeithel, Arne, 1998. "On the Ethnography of Cooperative Work," in Engeström and Middleton [eds.], *Cognition and Communication at Work*. New York: Cambridge University Press.

Resnick, Mitchel, 1994. *Turtles, Termites and Traffic Jams. Explorations in massively parallel microworlds*. Cambridge MA: MIT Press

Rogers, Yvonne and Judi Ellis, 1994. "Distributed Cognition: an alternative framework for analysing and explaining collaborative working," in Journal of Information Technology, 1994, vol 9(2), 119-128.

Rosch, E., C. Mervis, W. Gray, D. Johnson, P. Boyes-Braem, 1976. "Basic Objects in Natural Categories." *Cognitive Psychology* 8: 382-439

Saussure, Ferdinand de, 1965, *A Course in General Linguistics*, translated W. Baskin, New York: McGraw-Hill

Suchman, Lucy, 1998. "Constituting Shared Workspaces," in Engeström and Middleton [eds.], *Cognition and Communication at Work*. New York: Cambridge University Press.

Vygotsky, L. S., 1978. *Mind in Society*. Cambridge, MA: Harvard University Press,

Vygotsky, Lev Semenovich, 1987. *Thinking and Speaking* (translated N. Minnick). In Rieber and Carton [eds.] *The Collected Works of L.S. Vygotsky, Volume 1: Problems of General Psychology*. New York: Plenum Publishing.

Whorf, B.L., 1964. *Language, Thought and Reality: Selected Writings of Benjamin Lee Whorf*, J. Carroll [ed.]. Cambridge MA: MIT Press

Zager, David, 1999. "Mediating Heterogeneous Contexts with a virtual world." in *virtual worlds and Simulation Conference (virtual worldSIM'99)*, Landauer and Bellman (eds), Simulation Series v31.2 San Diego: Society for Computer Simulation.

7 Acknowledgements

Aspects of the virtual world described in this article have been constructed as the commercially available Trinity™ product from Avesta Technologies, which is the owner of any intellectual property associated with the product, whether or not implemented. Many thanks to Michael Brown of Avesta Technologies for his help with the editing of this article; to Bonnie Nardi of AT&T Research for her comments in an ongoing discussion of Activity Theory; to David Redmiles for his detailed commentary; and to the two anonymous reviewers who patiently waded through an earlier version and each made excellent suggestions about how to increase the signal-to-noise ratio of the paper.

8 Biography

David Zager, Ph.D., is Vice President and Chief Scientist at Avesta Technologies, Inc., which he joined shortly after its founding in 1996. As Head of Research, his responsibilities span technology strategy, architecture, and basic research. Before joining Avesta, he spent thirteen years at Morgan Stanley, where he was chief architect and lead developer of a cross-platform data infrastructure and middleware services research group. David came to software architecture after holding a postdoctoral fellowship in Cognitive Science at the University of Chicago in 1981-82, and teaching Linguistics at SUNY at Stony Brook. His ongoing research focuses on the nature of collaboration, particularly in computing operations.

9 Diagrams

Figure 1

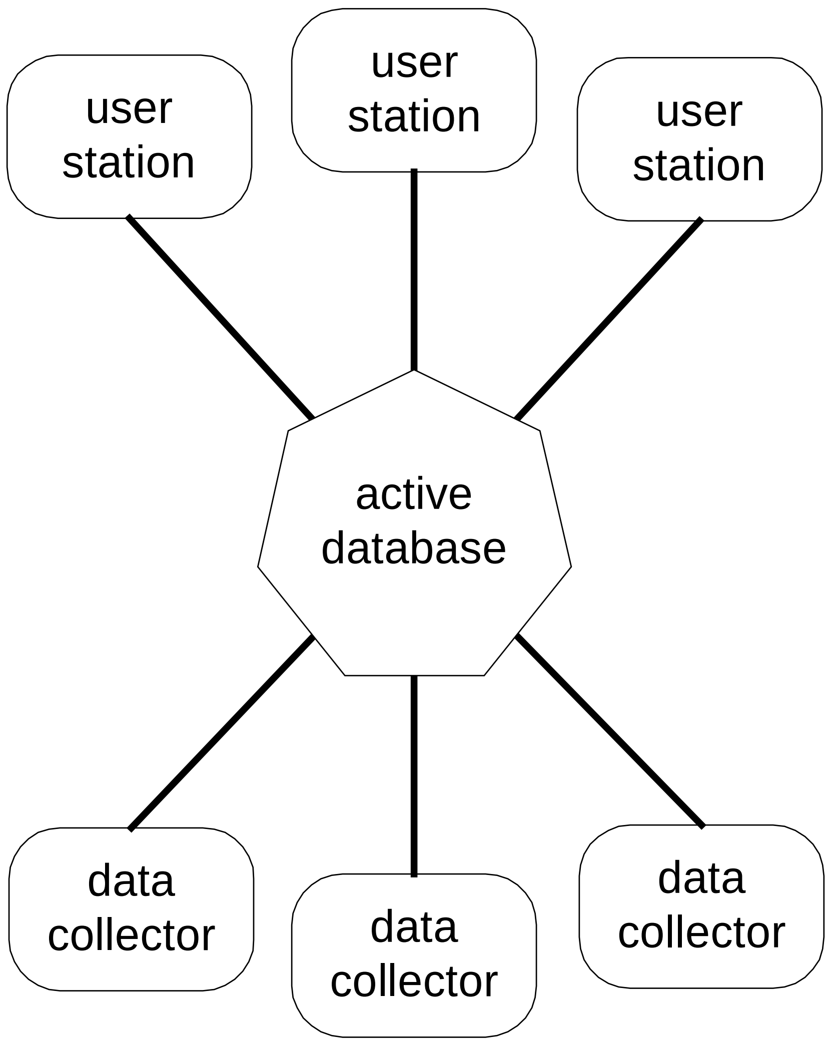

Figure 1: hourglass architecture of the virtual world

Figure 2

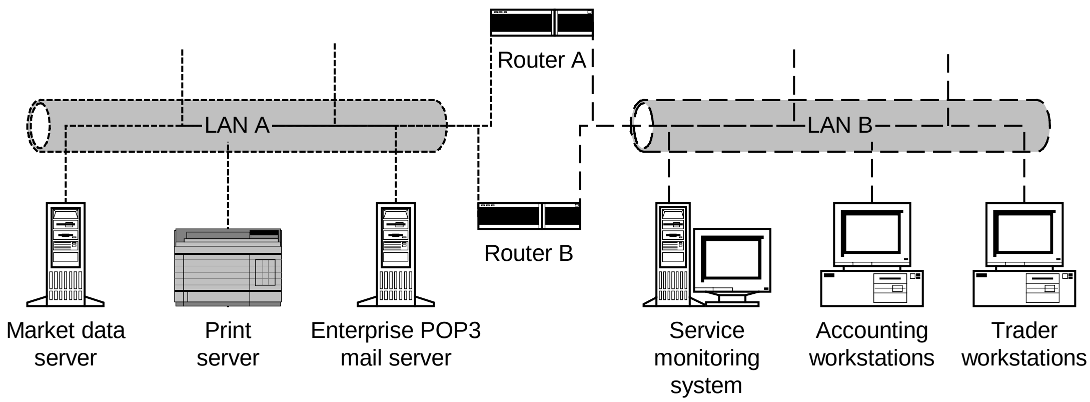

Figure 2: Example of a network fragment

Figure 3

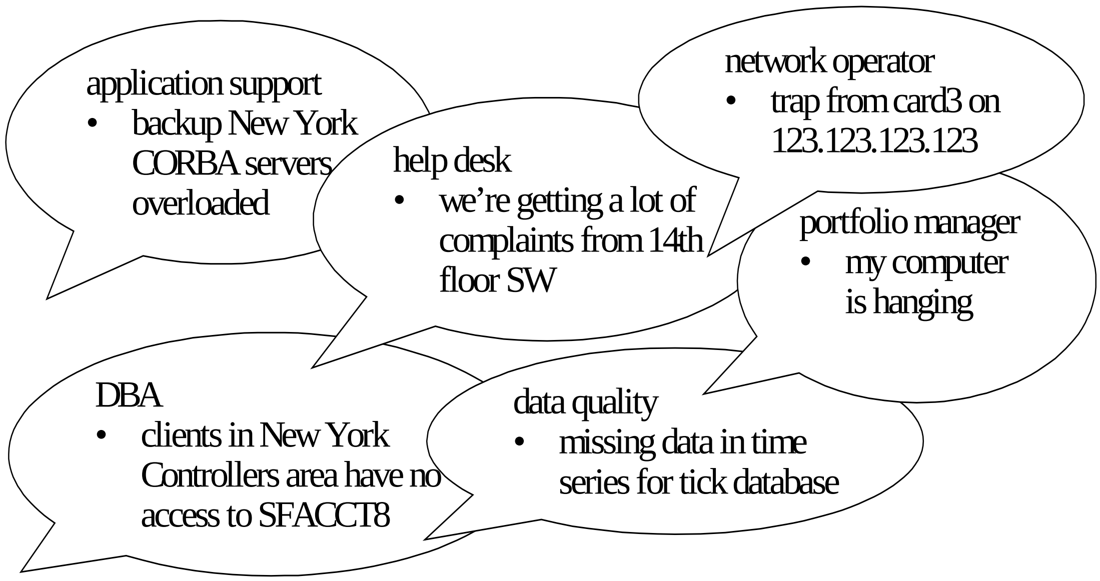

Figure 3: Different perceptions of a single service fault – association of apparently independent phenomena

10 Footnotes

[1] The context extends beyond the immediate group of participants, and stretches across the history of many different participants at different times and places, allowing the vocabulary and interpretation to evolve. But let us ignore the spatio-temporal aspects for the moment.