Massachusetts Institute of Technology
Sloan School of Management
Center for Coordination Science

A coordination theory approach to process description and redesign

Kevin Crowston and Charles Osborn

CCS WP #204 SWP # 4029

July 1998

77 Massachusetts Avenue, NE20-336, Cambridge, MA 02139
web: http://ccs.mit.edu/ccsmain.html

A coordination theory approach to process description and redesign

Kevin Crowston *

School of Information Studies
4–206 Centre for Science and Technology, Syracuse University
Syracuse, NY 13244–4100

crowston@syr.edu
+1 (315) 443–1676
FAX: +1 (315) 443–5806

Charles S. Osborn

Babson College
Wellesley, MA

osborn@babson.edu
+1 (617) 239-5585
FAX: +1 (617) 239-6416

Submitted to *Information and Management*

Category: Techniques

* Please address all correspondence to the first author.

A coordination theory approach to process description and redesign

Abstract

Managers must understand, influence, and redesign organizational processes to improve business performance. In this paper we present a technique for documenting a business process. The technique has six steps: defining process boundaries, collecting data, determining actors and resources, determining activities, determining dependencies and model verification. While similar to other process-mapping techniques, our approach is novel in incorporating ideas from coordination theory, thus the attention to dependencies. As a result, the technique is useful both for documenting a process and suggesting ways in which the process could be redesigned. We present an extended illustration with the hope that the technique can be used by readers of this article.

Keywords: coordination theory, process analysis, business process redesign, qualitative data collection

A coordination theory approach to process description and redesign

Introduction

Most managers develop usable understandings of the work they and their colleagues do, but the scope and complexity of their work practices often makes it difficult to comprehend them fully. This deficiency becomes most apparent when the way the work is done must be changed, for example, as information technologies are deployed to support or partially automate the work. Understanding is particularly difficult for semi-structured, knowledge-oriented work, where the flow of work is not reflected in physical production lines. Our goal is to help people understand their work as a prelude to changing and improving the way they do it. In this paper we propose and demonstrate a technique to analyze and represent work based on coordination theory [12].

We were trying to satisfy several potentially conflicting requirements in developing this technique.

- First, and perhaps most importantly, we wanted a technique that is generative, that is, capable of not only documenting what people do now, but also suggesting feasible alternatives.
- Second, it was important that documenting a process not become an end in itself. While understanding the current process is important, documentation is not the only aspect of organizational change that managers must consider, and probably not even the most critical one. Therefore, as a general principle, we add complexity to a description only if it helps answer some question important for a redesign.
- Third, we want the technique to be valid, in the sense that the suggestions of the technique must make sense to the individuals involved in the work.

To achieve these goals, we were willing to sacrifice a degree of reliability, in the sense that two analysts studying the same process will develop exactly the same description. Put alternately, we wanted our models to be simple and general, at the possible cost of accuracy. Instead, we apply a looser but perhaps more practical criterion: one analyst studying a process in some context will derive descriptions that can be readily understood and debated by another. Furthermore, the analysts should be able to combine their individual descriptions into a jointly acceptable representation that incorporates the characteristics identified by each. Such a representation might serve as the foundation of an explicit consensus between different analysts that recognizes a shared interpretation of the configuration and priorities of process details.

Our resulting technique is broadly similar to other process mapping techniques. We describe processes as *sequences of activities* performed by organizational *actors* that produce and consume *resources*. As well, drawing on coordination theory [12], we explicitly search for and represent *dependencies* within the process and *coordination mechanisms* used to managed those dependencies. Coordination theory has been used as the basis for a number of analyses [e.g., 2] but until now, there has been no description of how to apply the techniques.

In the next section, we review the theoretical bases for our technique, with particular attention to dependencies and coordination theory. In section 3 we walk through the stages in our proposed technique by presenting an extended example and in section 4, we discuss how the representations developed can be used to suggest alternative processes. We conclude by presenting an evaluation of our technique and its implications for action.

Theoretical basis: Processes, dependencies, and coordination

In this section, we review the theoretical bases for our technique, briefly discussing processes, coordination theory and dependencies.

Processes

In the past few years, "business process" has become a potent buzz-word for those interested in organizational change. Practitioners usually define "business processes" as sequences of goal-oriented actions undertaken by work units or business firms that repeat over time and which are measured in performance terms, such as time, resources expended or costs [e.g., 5; 9]. For example, Davenport and Short [6, p. 12] define business processes as "logically related tasks performed to achieve a defined business outcome." Harrington [9, p. 9] defines processes as, "any activity or group of activities that takes an input, adds value to it, and provides an output to an internal or external customer." In both definitions, key elements are activities, actors and resources.

In our work, we build on these definitions of process. However, we acknowledge that the relationship between work and its description can be more problematic. Except in the most routine processes, people do not do exactly the same things, and yet we want somehow to identify a set of "repeated activities" in the process they perform. Even identifying a particular set of things someone does as an "activity" can be difficult. It may be easy to label the one-minute cycle of an assembly line worker, for example, but finding the boundaries between coming up with an idea and writing it down (for example) is much more difficult. In general, though, we adopt a pragmatic attitude towards these issues. It is true that there are problems in representing processes, but in most cases it is possible to develop a meaningful and recognizable model. Therefore, our criteria in assessing a model is not some Platonic ideal, but rather that it makes sense to the users of the models, or at least, that users are able to come to some agreement about them. This stance is similar to Checkland [1], who describes models not as "what is" but rather as "what I shall (temporarily) take things to be in my analysis" (p. 175).

Furthermore, instead of arguing that our models are somehow true representations of work, we view the description as a discursive product, that is, as an artifact, with an author, intended to accomplish some goal. Checkland [1] similarly describes models as "opening up debate about change" rather than "what ought now to be done" (p. 178). Descriptions are resources for action, that is, someone doing the work may find them useful as a reference or justification for particular actions (this formulation was suggested by Lucy Suchman in an unpublished presentation at a University of Michigan CREW workshop on process modelling). Particularly important for this paper, someone may find a description useful as a basis for suggesting changes in the processes. Our goal in this paper is to describe how we build such potentially useful process models.

Coordination theory

A major drawback to many process representations is that they are, ironically, static. They describe the current state of a process more or less faithfully, but they do little to illuminate possible changes or improvements. We chose to use coordination theory as an approach to making generative models (i.e., models that suggest alternative ways a process could work).

As mentioned above, we analyze processes in terms of actors performing activities to achieve desired goals. According to coordination theory, these actors face *coordination problems* arising from dependencies that constrain how tasks can be performed. Coordination problems are managed by activities that implement *coordination methods*.

The first key claim of coordination theory is that dependencies and the mechanisms for managing them are general, that is, a given dependency and a mechanism to manage it will be found in a variety of organizational settings. For example, a common coordination problem is that certain activities require specialized skills, thus constraining which actors can work on them; this dependency between an activity and an actor arises in some form in nearly every organization. Coordination theory suggests identifying and studying common dependencies and their related coordination mechanisms across a wide variety of organizational settings.

The second claim is that there are often several coordination mechanisms that can be used to manage a dependency. For example, mechanisms to manage the dependency between an activity and an actor include (among others): (1) having a manager pick a subordinate to perform the task, (2) first-come-first-served and (3) a labour market. Again, the claim of coordination theory is that these mechanisms may be useful in a wide variety of organizational settings. Organizations with similar goals achieved using more-or-less the same set of activities will have to manage the same dependencies, but may choose different coordination mechanisms, thus resulting in different processes.

Taken together, these two claims suggests that alternative processes can be created by identifying the dependencies in the process and considering what alternative coordination methods could be used. Therefore, looking for dependencies and coordination methods is a useful start to process analysis and redesign.

Dependencies

Dependencies and coordination have been studied by many organizational researchers. Researchers have typically conceptualized dependencies as arising between actors rather than between the tasks the actors happen to be performing. The cause of a dependency is variously viewed as control by one actor over outcomes of actions of another or due to exchanges of resources [e.g., 15].

Although there have been methods proposed for measuring the strength of interdependencies between organizational entities, these approaches have tended to be unitary to a degree that forestalls their use for analyzing process behavior. As well, less has been written about how to identify dependencies within processes and how to apply dependency-oriented analysis to process redesign.

Crowston [4] categorized dependencies between activities by examining how the activities use common resources; this idea has since been greatly extended [13]. A tentative typology of dependencies is given in Table 1. Particularly prominent among these are *flow* dependencies, in which resources produced by one activity are consumed by one or more subsequent activities, and *sharing* or *shared resource* dependencies, where two or more activities share the same resource(s). Assembly lines, for example, offer mechanisms for managing the flow dependencies that exist between the activities required to assembly an automobile; budgeting processes offer an opportunity for multiple departments to negotiate the sharing of a fixed financial resource.

Insert Table 1 about here

Note from Table 1 that some dependencies decompose into subparts. For example, a flow dependency can be understood as including at least three subdependencies: a *usability* constraint that the resource produced by one activity is indeed appropriate for use by consuming activities, a *prerequisite* constraint that the producing activity needs to occur prior to the consuming activity, and a *transfer* or *accessibility* constraint that the resource is readily available to the consuming activity. Each of these constraints needs to be accounted for in analyzing the overall flow dependency. As well, most processes can be understood as including a mix of dependencies. Assembly lines that are timed to the capacity and speed of their most expensive machine manage both flows from activity to activity and the machine capacity (a shared resource).

Not surprisingly, knowledge-intensive work is often coordination-intensive. Knowledge workers within an organizational hierarchy are often asked to adjudicate conflicting claims on resources in order to maintain acceptable levels of process performance. For example, production schedulers coordinate assembly lines in order to maximize asset utilization and output flexibility. In this context, coordination roles that cross organizational boundaries often represent a response to dependencies that affect important parameters of organizational performance.

For example, consider a company where an account executive (AE) develops job quotations for customers and is also responsible for supervising the quality of internal work required to complete that job. In this role, the AE coordinates a process that crosses the boundary between the company and its customers—managing a flow dependency that critically affects the usability of the company's output to its customer as well as the usability of the customer's input (e.g., the quote) to the company. It is not difficult to see how the success of the AE's organization depends greatly on the success with which the AE manages this dependency. In this paper, we will consider an example of such as cross-boundary dependency in some depth to illustrate our process-analysis technique.

A coordination theory approach to processes description

In this section, we describe our process-description technique in six stages. We start by setting the boundaries of the process to analyze. Second, sources of data on the process are identified and data collected. The heart of the analysis is the identification of activities, actors and resources and dependencies between them (steps 3, 4 and 5). Finally, the process model must be verified. Although we present the steps in this order, in practice, analysis and data collection are likely to be interleaved, as analysis reveals gaps in understanding, which motivate further data collection. For example, some data are necessary to set the process boundaries, while the process boundaries are necessary to bound the data collection. Similarly, the steps in an analysis will be performed iteratively, as a greater understanding of one aspect of the process will suggest additional alternatives to consider in the others.

With the exception of step 5, our technique is similar to many others. While we will present the approach we use in our own work for the remaining steps, we do not claim that they are the only way to gather and analyze data (or even always the best). If you prefer the approach of some other process modeling technique, then those techniques can also be used, as long as they identify goals, actors, resources and activities for use in step 5. In particular, coordination theory might be a useful conceptual model to use as part of a soft-systems analysis [1] as the root definition of the relevant systems includes the necessary elements (p. 164; see also the discussion of CATWOE, pp. 224–227).

Source of examples

We will illustrate our technique using examples drawn from a case study of a small marketing services company. As Checkland [1, p. 192] notes, "Authors would better keep their models and methodologies to themselves until they can demonstrate a problem solved by the use of them." In that spirit, we offer this case as a demonstration that our method has helped at least one organization. We chose this company because its core processes are simple enough to present in a paper, yet sufficiently complex to permit a discussion of generative process design. To motivate our discussion, we will first briefly describe the company, which we will refer to by the pseudonym "MAG Services" or simply "MAG".

MAG Services is a wholly owned subsidiary of a direct mail company that provides mailing and inquiry fulfillment services for corporate marketing departments of Fortune 500 corporations. MAG receives requests for information about a client's product(s) from individuals and fulfills the requests by mailing out appropriate marketing materials. MAG provides two kinds of service: custom and non-custom. Figure 1 shows the basic work flow through MAG's facility for both kinds of job.

Insert Figure 1 about here

- In the *non-custom business*, fulfillment is similar for all clients. A typical job would work as follows. A company runs an ad in *Business Week*. The ad includes a tear-off postcard bound into it saying, "Send me more information about…" and offering choices of products (e.g., "I want to know about blue widgets, or large widgets, or oil-resistant widgets", etc.). On the front of the card is MAG's address. A filled-out card with a return address, demographic information, and product interest arrives by mail at MAG's data-entry room. MAG mails back product brochures according to the selections made by the person who sent in the card, and collects the data from the cards to feed back to MAG's client as marketing leads. In non-custom work, jobs typically run for an extended period (usually longer than a year); mailing materials are relatively standard and supplied in bulk by the client; mailing and production tasks do not vary significantly over time; and MAG performance reporting is largely limited to tracking the number of qualified sales leads.

- *Custom business*, on the other hand, is performed for clients on a one-time basis. An example of this non-traditional work might be a contract in which MAG provides inquiry fulfillment services following a trade show. At the show, anyone passing by the booth can tear off a postcard with "Send me more information…" on the back of it. Fulfillment of these requests is similar to the non-custom case. Based on these inquiries and other data, MAG creates a database of leads, including the addresses of people who have inquired, the inquirers' demographics, lists of pre-qualified leads, structured customer feedback, and the like. In contrast to non-custom jobs, though, these jobs run for only a limited time with a concentrated volume of work, materials may be specific to the particular job and include materials customized to the requester, and performance reporting includes both volume of contacts and measures of the quality of leads generated.

During a ten-year period in which MAG became a leader in the traditional mailing services business, the company developed a sophisticated database system that produced both most of the mailings required by high-volume, standardized inquiry fulfillment and the reports required by long-time clients. By 1994, the company was actively engaged in applying the same database to customized services. By that year customized services accounted for more than 40% of revenue. However, at the same time MAG's management recognized severe operational and profitability problems with the custom business. These issues captured management attention because custom contracts were straining the capacity of the organization and customers dissatisfied with MAG's performance on custom work were beginning to direct follow-on business to competitors. MAG managers gave high priority to custom work because they believed that it represented the area into which the company would have to grow in order to maintain its market share within the mailing services industry.

In the remainder of this section, we will describe how we analyze the work done in MAG to suggest alternative processes that might be more efficient or effective, working through our technique step-by-step.

Step 1: Setting process boundaries

The first step in our analysis is setting the process boundaries. Boundary-setting involves decisions about which actors, resources, and activities are central to the analysis and which are included only as tangential links to other processes or not at all. Checkland [1] notes that "there is in principle no limit to the analyst's freedom to make whatever choice he thinks or feels might lead to insight" (p. 221).

We choose to focus our analyses around the stated goal of the process. Activities, actors, and resources that contribute to this goal are included in the analysis; activities that are peripheral are included only abstractly or not at all. In many cases there may be multiple possibilities for the overall goal of the process. In these cases, the process boundaries are particularly important because they define the "problem" and thus the scope of the solutions considered [18].

Identifying goals

An obvious way to identify goals is to ask process actors why they perform process activities. However, Spradley [19] is quite emphatic about not asking people "why" questions. He points out that intentions and motivations are subject to a great deal of rationalization as well as interviewer "demand characteristics", so much so that their reported intentions must often be handled with suspicion. Instead of asking "why", Spradley suggests asking, "Under what circumstances would you do X?" or "Describe a situation where X would be appropriate." This may be a little too rigid for all situations, but we believe it is worthwhile to keep in mind that there are always alternatives to simply asking "why".

For the top level goals, we ask more general questions about the business purpose and overall objectives (e.g., "How does this process fit into the business?"). For the lower level activities/goals, we suggest more specific questions, such as:

1. What purpose does this activity serve? If you stopped doing it, what would happen?
2. How is (or how could) performance of this activity measured? What counts as a "good" or "bad" performance?
3. Who uses the results of this activity? What kinds of results to they find most important or helpful?

For MAG Service's custom business, these questions produced the following answers:

1. At the highest level of abstraction, MAG Services and its customers had divergent *purposes* in entering into custom business arrangements. Customers wanted to find new customers through trade shows and targeted mailings, rather than waiting for customer to come to them as before.

MAG wanted to grow sales and profits by generating new jobs that deliver such contacts with the promise of follow-on business. Custom-designed mailing programs were intended to satisfy both these goals.

2. *Performance measurements* reflected the potential divergence of these goals. MAG's customers were interested in new sales based on leads generated by customized mailing services. MAG was interested in the profits from custom work and follow-on business.

3. The definition of *helpful results* differed in the same manner. Customers wanted to learn more about their markets, so as to build sales. MAG wanted to learn more about what services customers needed, but in a way that enabled the company to deliver those services at reasonable cost.

In other words, our initial study of the process suggests that at a general level this process has two divergent goals, first concerning customers' needs for specific attention to generate an explicit sets of sales leads and second concerning MAG's objectives of assembling and delivering mailings as efficiently and profitably as possible. Each of these views—customer and company—may be appropriate depending on the purpose of the analysis. The point is not whether there is a single, universal process perspective that fits all conditions, which seems unlikely; instead, the more appropriate question to ask is whether the boundary of the process chosen is appropriate for the problem the analysis is intended to address. Subsequent analyses can demonstrate how the process contributes to resolving rather than reinforcing divergent goals.

Furthermore, as an analysis proceeds, it may be useful to change the definition of the process boundaries under consideration. It is not a question of the definition being right or wrong as much as useful or not useful. As Checkland [1] puts it, "the systems thinker must be able cheerfully to abandon his earlier choice of relevant systems and start again" (p. 223), perhaps shifting focus from a stated primary task to some latent issue that must continually be addressed (p. 222) or moving higher in a process decomposition hierarchy. For example, analysts considering supply chains might start with a process representation that describes one participant in the chain, but expand the analysis might to include multiple flows coordinated by multiple corporate actors. Checkland [1] argues strongly that the definition of one system must take into account the definition of the other systems it is to serve (pp. 234–5). For example, MAG might develop different strategic goals if it saw itself as an integral part of its customer's marketing processes rather than as a provider of standardized marketing services.

In our example, the focus of the analysis will be on the company that provides mailing services because the management problem under consideration is how to position the company for profitable growth. Other questions might lead to different boundaries (i.e., the definition of the system under study depends on the purpose of the study rather than being an inherent property of the system). For example, a study of companies that used account executives as sales coordinators might contrast MAG with several of its competitors using higher-level process maps.

Step 2: Collecting data

Building a process representation requires collecting considerable detail about many activities, goals, actors and available resources, as described in the following sections. In this section, we describe our approach to data collection. Many data collection techniques have been proposed, which make different tradeoffs among rigor, speed, cost, and accuracy. Many of the techniques focus on the question of reliability, in the sense that a second observer using these techniques should come to the same exact conclusions about the organization. Such rigor is clearly necessary for doing scientific studies where the goal is to make some generalizable assertion about how some phenomena works in multiple settings. However, our goal here is different as we simply want to say something about a particular site that others will find interesting or useful. In particular, while we argue that the same coordination mechanisms are found in many settings, we do not require that analysts will necessarily agree on the mechanisms used in a particular site (although one should at least recognize the set proposed by another analyst).

In our own work, we draw on four field-based research traditions: grounded theory [7], ethnographic research [19], case study research [21] and the clinical perspective in fieldwork [16]. We have found these perspectives useful because they focus on the categories and terms that process participants themselves use to describe the process while imposing a minimum level of external preconceptions on process representation.

The methods we use include three central components:

1. Semi-structured interviews based on understanding process decompositions, specializations, and dependencies;
2. Observation and participant observation where such approaches appear appropriate (this may include a range of participation from "stapling oneself to an order" to sitting in on meetings).
3. Iteration that encourages revisiting collected data repeatedly as process understanding grows during subsequent phases of analysis.

As well, we examine existing data about the process, such as flowcharts of processes and process fragments, examples of documents created in the process, training manual or even interviews with managers or narratives collected from line workers. Even if this evidence was originally collected for purposes other than process analysis, it can be used to increase understanding.

At MAG Services, we used a combination of these steps. One of the authors interviewed more than 15 members of an organization of 70 people, including all of the company's account managers, the managers comprising the top three levels of the company, and selected part-time employees (e.g., in data entry and operations). Participant observation accumulated over approximately 20 weeks over eighteen months, and was largely accomplished by a MAG Services manager whom we trained in process analysis techniques.

Step 3: Identifying actors and resources

As data is collected, we begin to create and populate the various categories of our description. Although we have presented this as a distinct step, note that in practice data collection and analysis are likely to be interleaved (i.e., Steps 3–5 and Step 2 overlap). In step 3, we identify the actors who execute the process and the resources used and created in the process. This step is also useful in refining the process boundaries, because only activities performed by the selected actors around the selected resources will be included.

Identifying process actors who are direct human participants is relatively straightforward. Non-human actors are more difficult to identify. In some cases machines might be viewed as actors (e.g., the "the database sorted the leads by zip code"); in others larger aggregations, such as departments, might be considered ("Sales qualified the customer"). Again, our general rule of thumb is to add detail only where necessary for the purposes of the study. For activities on or near the boundaries of the process description, for example, aggregations may be appropriate for defining actors ("Federal Express takes the package from the loading dock…"), while for activities central to the process description, human and system-related actors may need to be described more specifically ("Martha reads the output from the quality testing equipment on the assembly line. She's the only one of us who knows how to interpret it, and she's almost always right"). Similarly, we would tend to treat a computer system as an actor unless we there were some reason to concern ourselves with the source of data or the programs embedded in the system. However, it is worth noting Checkland's [1] advice to "avoid the unthinking assumption that organizations, departments, and sections are *ipso facto* systems" (p. 226).

To check that the set of actors is complete, we follow the work flow up and downstream, using questions like, "After you get done, to whom does this paperwork go?" or "From whom do you get your work?". The tracing can diminish (and the actors become increasingly aggregate) at the edges of a selected process boundary.

Once the set of actors is identified, we group them into classes of actors who perform similar activities in a similar fashion (i.e., who fill similar roles in the process). For example, we might chose to treat all accounting clerks or FedEx drivers as examples of a class and document how these actors work in general as opposed to in particular. A possible rule of thumb for this grouping is that a given task could be performed by any member of the group. At MAG Services we developed the list of actors shown in Table 2.

Insert Table 2 about here

As the actors are being identified, we also begin to list the resources that are created by or pass between activities. Some of these may be physical objects, which are relatively easy to identify. Often, however, the key resource is information. To identify information resources, we ask what messages the actors send one another. For example, at MAG, the account executives (AEs) write instructions for all the different production participants, e.g., data entry instructions, technical work orders (for programmers and Technical Services workers), and mailing services work orders (e.g., operations/production people in the warehouse). Resources identified in the case are shown in Table 3.

Insert Table 3 about here

Step 4: Identifying activities

The next stage in our analysis is to list the activities that compose the process (i.e., what gets done in the process). Several problems must be addressed in this analysis. First, the same activities may be labeled and interpreted differently by different actors or vice versa. Second, activities can be described at varying levels of detail. Third, the actual activities observed may vary between performances of the process, making identification of "the" process problematic. The solutions chosen for these problems have implications for understanding activities themselves and for identifying the resources that activities use.

How to identify activities

Spradley describes an interviewing technique used by cognitive anthropologists in collecting complex information from informants [19, esp. chapters 4-6]. The basic technique comes in three parts.

1. One can ask "grand tour" or "mini-tour" questions. These are general, open ended questions like, "Tell me about a typical day. What do you do?" A grand tour potentially covers all of the different work activities that go on in a particular task unit.

2. One can focus on more specific topics, like a particular part of the job or a particular process: "Tell me what happens when a customer comes in" and so on. One can follow up, probe, ask for elaboration, clarifications, etc., but the basic idea is to elicit a general description of the work in the respondent's own words.

3. Having obtained a basic outline for groups of activities, it is possible to focus the contents of various "semantic domains" [19, p. 107]. A semantic domain is like a category, which can be very high level or very low level. In understanding process decomposition, "activities", which can be decomposed into "steps in" or "parts of" the activity are important. To get at this information, one can ask a series of structural questions: "What are all the steps in this process?", which is followed up with various probes to check for completeness: "So far, you've mentioned the following steps: a, b, c, d, e,… Are there others?" At the most detailed level, it is worth checking that all activities are direct, meaning that they are something an actor can actually do (e.g., collect information or make plans), as opposed to the hoped-for outcomes of an action (e.g., lower costs) [1, p. 235].

Decomposition of a process into activities

Given a description of an activity, a second problem is to choose the appropriate level of decomposition at which to represent it. For example, a flowchart might include one box for a particular actor's task or hundreds of boxes for the fine details of that task. Each level of detail might be appropriate for different purposes. We avoid this problem by developing process decomposition trees that simultaneously represent varying levels of detail. (Many dataflow diagram techniques offer the same advantage.) This approach also makes it possible to view a process at a consistent level of detail by picking the same level of the process decomposition trees.

For example, at MAG Services, we note that "Send Mailing" is an atomic activity, but that "Run Job" can be broken down into at least four sub-activities, including preparing quotes, setting up jobs, producing jobs, and providing status reports. Following our general rule of thumb, we keep decompositions at the most general possible level unless the problem to be analyzed provides a reason to decompose a process in more detail.

Variations on a process

A third problem is representing variations on the process. Coordination theory starts from the premise that many different sequences are possible, and that these will differ from time to time. In many processes, the exact activities observed in a single instance of a process may never be repeated in all particulars. Even so, most organizational participants have little difficulty in recognizing the process as an abstract description that represents multiple instances of specific steps, both those that have happened and those that might happen in the future. This feature of organizational life can be described as inducing a generalized process from a relatively small set of observed activities by fitting observed actions into mental templates that define more abstract process steps. This translation usually forms an important part of "learning the ropes", the acculturation that orients newcomers to existing organizational characteristics and habits.

As well, the abstraction will suppress particular activities. For example, two people may start each interaction by spending a few minutes discussing last night's game, yet not include that activity when describing the process. In general, we would follow the informants' lead in choosing whether to include such activities in the process description. It may be that these interactions are viewed by some of the individuals as necessary to the smooth running of the process, in which case they need to be included; or it may be that they are considered as secondary, in which case we would probably also leave them out.

Mini-tour questions can generate at least a reasonable first approximation of activities that occur in a process. It is trickier to get a definitive sense of activity sequence and the contingencies that may result in changes in the sequence. Getting a respondent to describe "the" sequence is unrealistic, yet we wish to represent typical sequences, in a way that is recognizable and acceptable by participants.

Our approach to this issue is to build a prototype sequence and then asking for variations. For example, ask the informant to describe a typical sequence of steps. Then, ask: "Can you think of an example where the steps were done is a different order?" or "Under what circumstances would you do things differently?" Equally important would be examples of abortive or failed sequences: "Is there ever a time when you stop in the middle?" Comparing sequences and sequence variation provide useful data for identifying and understanding dependencies, which are discussed further below.

Example

A decomposition of the activities in the MAG Services example are shown in Figure 2. This figure shows a hierarchical decomposition of the process of providing MAG services, shown at the top of the page, into activities and subactivities, drawn down the page. This description shows that MAG engages in three phases of activity when handling a typical job. These phases include qualifying prospective customers, providing custom and non-custom mailing services, and billing clients.

Insert Figure 2 about here

The activity of providing mailing services is itself further decomposed in Figure 2. The operational details of providing mailing services are normally handled by MAG account executives (AEs). The AE writes and distributes several sets of instructions inside the company so that the mailing which is produced ultimately matches the client's specifications (as approved in the quote). The AE also stays in touch with the client to ensure that MAG's services continue to be satisfactory as the job progresses. Most of the examples discussed below concern coordination managed by AEs in initiating new work on behalf of the company.

Figure 2 shows that MAG undertakes mailing contracts by qualifying prospective customers, providing services, and billing clients. Once prospects are qualified, the company sells its services. The sales activity ends when an AE takes an order for a job expressed within the company as a job quotation. Once the quote is approved, the AE prepares and distributes instructions that describe the job to operational departments within the company, and shows a copy of sample output from the job to the client. Finally, the operational departments execute the instructions and complete the job, which results in the collection of market information and a completed mailing. Sales and billing functions are performed in part by MAG's parent company.

Earlier we discussed an important variation in the processes that MAG Services used to deliver mailings: the difference between custom and non-custom work. Figure 3 suggests how this variation can be interpreted as representing variations in the Provide Services process. Figure 3 shows that the generic process of providing mailing services can be provided in to different ways, indicated by the two different specialized forms of the process—providing custom services and providing traditional services—drawn below on the page. Note that the specialized versions of the process have specialized versions of the subactivities as well. This comparison enables us to focus in useful ways upon the differences in producing mailings for custom and non-custom work. Table 4 summarizes some of these dimensions from the point of view of MAG's management.

Insert Figure 3 and Table 4 about here

This comparison surfaces some of the coordination challenges raised by custom jobs that a focus on production efficiency alone would not recognize. Non-custom work is long-term, standardized, brings guaranteed mailing volumes, and operates against a schedule largely set by MAG. Custom work is short-term, non-standard (even with respect to the order in which specific production tasks are done), operates against tight deadlines, and requires daily contact with the customer. Custom work differs fundamentally from traditional jobs in process and coordination.

Step 5: Identifying dependencies

So far, our technique resemble most other process mapping techniques, identifying activities, actors, and the flow of resources. These results for our example are summarized in Table 5. The novel aspect of our approach is the identification of dependencies between the activities and resources and the application of coordination theory, which we discuss in this section.

Insert Table 5 about here

Given a process description that includes goals, activities, actors, and resources, we propose two general heuristics for identifying dependencies.

1. *Dependency-focused analysis.* Identify dependencies, then search for coordination mechanisms. In other words, look for dependencies, then ask which activities manage those dependencies. Failure to find such activities might suggest potentially problematic unmanaged dependencies.

2. *Activity-focused analysis.* Identify coordination mechanisms, then search for dependencies. In other words, identify activities in the process that appear to be coordination activities, then ask what dependencies those activities manage. This approach asks directly whether all observed coordination activities are necessary.

These approaches are described and illustrated in the remainder of this section.

Dependency-focused analysis

In dependency-focused analysis, we examine the activities and the resources they use, determine possible dependencies by considering which resources are used by more than one activity and then look for other activities within the process that manage these dependencies. More specifically, to identify dependencies and mechanisms, we ask questions such as the following about each activity in turn:

- What are the inputs to this activity (physical, informational and other necessary preconditions, such as permissions)? Are there flow dependency with the activities that create these resources? Are these resources used by other activities, creating shared resource dependencies?
- What are the outputs? Is there a flow dependency with the activities that use these resources? Do multiple activities create these resources, creating common output dependencies?
- What other resources are used, e.g., actors, equipment, overhead, time, or other items of importance in the process? Are there shared resource dependencies with these resources? How are these resources assigned to this activity?

- What performance problems have been reported for this process (e.g., observed divergence from stated goals)? Do these problems reflect unmanaged dependencies?

For each potential dependency identified this way, we then search for activities that manage it. The typology in Table 1 is helpful, as it suggests a range of possible coordination mechanisms for each type of dependency. For example, if an activity needs a resource, then from the typology we note that the resource may be permanently assigned, taken first-come-first served from a pool of resources, assigned by a manager, etc. A flow dependency might be managed by a single activity or the coordination mechanism might be decomposed into separate activities for managing the transfer, usability and inventory dependencies.

To summarize, in dependency-focused analysis, we examine the use of resources in order to identify potential dependencies, and then look for activities that manage those dependencies.

Example of dependency-focused analysis

The analysis discussed above can be done at every level of decomposition. We will illustrate by first considering dependencies in the MAG case at a very abstract level, considering the company as a link in a value chain, as shown in Figure 4. At this level, there are several resources, such as mailings and market information, that are used by multiple activities, thus creating dependencies. The dependencies between the activities are indicated using curved lines to show the follow of resources from one activity to the next. More specifically, examining inputs and outputs suggests that MAG produces resources for the "Using Market Data" activity (i.e., it provides inputs to this activity), including sales leads generated by MAG mailings and market information collected from sales inquiry forms (e.g., demographics, channel sensitivity, etc.). The "Using Market Data" activity is part of some larger process, indicated schematically by the vertical lines that connect this activity to other, unshown, activities.

Insert Figure 4 about here

Taken together, these resource uses suggest two kinds of dependencies: first, a task-resource dependency between MAG's customers and MAG, shown by the fact that MAG performs certain activities on behalf of their customers and second, a flow dependency between the activities of MAG and its customer, as shown the by flow of sales leads from MAG to the customer. We next attempt to identify the activities that manage these potential dependencies.

MAG's business starts when a customer decides to hire them to provide mailing services that leads to useful demographic data or qualified sales leads. We note that numerous activities, such as mailing marketing information, are performed by MAG on behalf of a client. Such an assignment of tasks suggests a possible task-actor dependency (a special case of a task-resource dependency). In other words, a customer needs these services, but does not or can not perform them themselves and therefore, decides to hire MAG to perform them.

The various activities needed to manage a task-actor dependency are shown in Table 1. These include determining needs, identifying possible actors, collecting information, picking the best and then assigning the task. Interestingly, in this case we see the assignment from the perspective of the assigned company, as it responds to requests for information ("Sell to customer" and "Prepare quote"), is assigned the job and actually performs it. As well, several other activities, such as "Qualify prospects" and "Bill client" are likely involved in managing this dependency, although in this case, these activities were performed by MAG's parent organization.

The business completes a service cycle once customers receive data in a manner that disposes them to seek more work from MAG (e.g., the "Use market data and sales leads" process in Figure 4). The service cycle is a process that manages a flow dependency existing between MAG's activities and those of their customers.

Coordination theory suggests that a flow dependency includes usability, prerequisite, and transfer constraints that influence process performance. Following this distinction, we can identify activities or groups of activities within the "Provide MAG Service" process that manage such constraints (see Figure 4).

Figure 5 provides a full overview of the process representation created so far in our analysis. It includes a hierarchical process decomposition, as in Figure 2, over laid with dependencies and coordinating activities, as in Figure 4. In Figure 5, we look within the "Provide MAG Services" process to understand how subactivities manage the dependencies that act as constraints on the flow of jobs. From the customer's point of view, key variables associated with process performance appear to lie within the "Run Job" process. For example, the time dimension of MAG's performance appears to be constrained by the speed with which MAG can set up jobs and produce mailings. This implies that "Set Up Job" is managing prerequisite constraints associated with the higher-level flow of jobs. The geographic nature of its work (e.g., disseminating mailings to inquirers and market data to clients) emerges clearly from the "Produce Mailing Service" process. This implies that "Produce Mailing Service" is managing transfer constraints associated with the flow of jobs. Preparing a quote is a critical step in ensuring that the job defined to the company is a job that will be satisfactory to the customer, which implies that "Prepare Quote" is managing usability constraints.

Insert Figure 5 about here

Figure 5 summarizes this analysis using a graphical notation that shows subactivities and subdependencies. The upper levels of the process representation describe MAG Services as managing a flow of resources between two of its client's processes, as discussed above. The darker arrows in the figure suggest how dependency-focused analysis moves downward within the activity hierarchy to identify coordination processes that manage subdependencies. In this case, specific coordination activities manage resources associated with a subdependency that constrains the flow of jobs. "Prepare Quote" manages the *usability* of a job to a customer and to MAG. "Set Up Job" ensures that the company completes the right tasks in the right sequence, thereby managing *prerequisite* constraints that affect the flow of jobs. "Produce Mailing Services" generates the physical mailings that fulfill inquiries and *transfer* information back to the client. These coordination activities are summarized below in Table 6.

Insert Table 6 about here

Using Figure 5 and Table 6, we can ask how effective the chosen coordination strategies have been in practice. Table 7 shows the results of a coordination analysis that explicitly considers coordination strategies. It compares the effectiveness of the coordination strategies that MAG developed for non-custom business with performance observed for custom jobs.

Consider the contribution that the analysis in Table 7 offers for understanding MAG's current business position. Recall that the company's existing coordination processes were designed for long-cycle, high-volume, low-variation jobs. Using our approach, it becomes possible to identify specific ways in which MAG's services are breaking down under the differing requirements of custom work. Specifically, custom business varies across dimensions such as deadlines, job complexity, and accuracy requirements in ways that MAG's existing coordination techniques are not particularly well prepared to handle. To illustrate, we will discuss three examples from Table 7 in more detail. The analysis enables us to apply the notions of usability, prerequisites, and transfers to specific operations-level activities within the company. These activities represent the coordinating mechanisms that the company uses, implicitly or explicitly, to implement its services. By this means we were able to pinpoint with some accuracy how coordination breaks down within the daily work practices of the company.

Insert Table 7 about here

Quotes: Coordinating usability. MAG's traditional work was sufficiently standardized that AEs could successfully negotiate quotes over the telephone, taking handwritten notes that were later revised into a quotation letter signed by the customer. This approach worked well for relatively simple standard work. Custom work, however, often varies in the types of services that customers requested, and always required much tighter deadlines. Under such circumstances, AEs did not always know how to quote jobs immediately, customers often didn't realize the cost implications of what they were asking for, and quotation letters became both delayed and increasingly controversial. In this sense, a quotation process developed for standard work proved unsuitable for coordinating quotes for custom jobs.

Job set up: Coordinating prerequisites. Once a quote was complete, AEs prepared and circulated instructions for entering data and producing a job. These instructions, delivered on internal forms that MAG designed for traditional work, became increasingly dysfunctional for custom jobs. The forms were long and complex; as the custom business evolved, their options became irrelevant to the instructions that AEs needed to provide. MAG's organizational systems, in effect, were asking for the wrong data. Some AEs reacted to this problem by taking more time to type their own versions of instructions; others hand wrote long additions to standard company forms. Others insisted on following up all written instructions with verbal instructions. The net effect of these reactions was to slow the pace of custom work at the very time that custom jobs were requiring faster turnaround times.

Mailing services: Coordinating transfers. MAG's internal operations were highly developed for producing standardized bulk mailings. Problems developed, however, when custom jobs required below-average batch sizes and MAG was unable to adjust. This problem surfaced when customers required AEs to report back to them on misdeliveries immediately rather than monthly or quarterly. Because MAG's reporting systems, developed for standard jobs, only traced activity by batch number, AEs had to spend hours researching potential mistakes. The net effect, again, was to slow down custom work and make MAG appear inflexible.

To summarize, in dependency-focused analysis we first identify dependencies by considering resources used or created by multiple activities. We then search for coordination mechanisms that manage those dependencies, searching through successively more detailed layers of the process until insights are gained about how process goals are implemented in practice. In a full analysis, this dependency focus leads to a detailed understanding of activities that coordinate key resources associated with dependency constraints.

Activity-focused analysis

Our second approach to finding dependencies and related coordination mechanisms starts from the activities. Activity-focused analysis surfaces candidate coordination activities, then looks for the dependencies that they manage. In this sense it operates inductively rather than deductively, aggregating dependencies upward through the process hierarchy to build an analysis that complements dependency-focused approaches.

In activity-focused analysis we suggest three complementary heuristics to triangulate on potentially important dependencies. These include identifying critical process tasks, identifying coordination activities, and identifying coordinators.

1. *Search for process-critical activities.* Activity-focused analysis asks which activities play a necesary role in the completion of a process; the remaining activities are likely to be coordinating these. At MAG Services, producing a mailing is a process-critical activity because it directly leads to the output desired by the customer.

2. *Search directly for coordination activities.* Activity-focused analysis examines tasks identified within a decomposition hierarchy and asks whether these activities represent coordination, i.e., whether they match one of the activities in Table 1 or otherwise manage an important dependency within the process. For example, examining budget preparation cycles can identify resource allocation mechanisms; tracking the flow of paper or other physical resources within an organization can often identify activities that manage flow dependencies.

3. *Search for actors or resources that coordinate.* Finally, activity-focused analysis looks for actors whose work frequently suggests coordination tasks. At MAG, account executives negotiate a contract and write the instructions that define customized mailing services.

To summarize, in the activity-focused analysis, we look for activities that may implement coordination mechanisms. Candidate activities are those that are non-production, resemble coordination mechanisms or are performed by coordinators.

Example

Figure 6 summarizes a search for coordinating activities at MAG Services. The illustration represents the results of the steps described above.

Insert Figure 6 about here

Search for critical activities. This step asks the same question of each activity: could the end product of the process exist without it? Of all the activities shown in Figure 6, the only one that appears irreplaceable is "Produce Materials". The company might use different sales processes; it can change its quoting process; it can even "send" mailings via the Internet; but jobs cannot be delivered without some production of personalized materials, be they physical, paper, or electronic.

This focus is helpful because it offers a core from which to aggregate dependencies. If "Produce Materials" represents a key production activity in this process, the analysis can step outwards from that foundation to ask what coordination other observed activities provide. From this perspective, the activities related to converting client interest into instructions for producing materials (9 of the 18 activities in Figure 6) appear to be attempts to coordinate production of mailings according to criteria that meet the performance expectations of both MAG and its customers (e.g., maximum profit with minimum time, errors, and cost).

Search for coordination. Within the "Run Job" process, the only obvious production step is "Produce Mailing Service", that is, only this step produces an output that is given to the customer. "Prepare Quote" and "Set Up Job" appear to be processes that largely prepare information to ensure either that jobs meet performance criteria (e.g., "Prepare Quote") or that work will proceed error-free

("Set Up Job"). In other words, they appear to manage the usability of the production work found in the "Produce Mailing Service" step.

Where records represent information that crosses process boundaries (i.e., they are an output that is used as input by another process), they identify potentially important flows. In this sense, records can form a resource within flow dependencies. By this means information-intensive activities (i.e., those that handle information used extensively by other activities) can often be understood as coordination mechanisms. Figure 6 shows records that cross the branches of the process tree developed for the mailing company. It describes six steps in executing a job: prospective customers produce an inquiry about mailing services, sales processes pass prospects and job proposals to MAG account executives, AEs prepare quotes in response to those proposals, quotes are converted into instructions, and instructions precede the mailings and leads generated by a job. Each of these elements represents a resource that flows across process boundaries within MAG's operations. In Figure 6 the "Run Job" process is shown at a lower level of decomposition because it represents internal activities over which MAG managers have greatest control (as noted above, sales and billing are performed by MAG's parent organization).

Search for coordinators. Actors perform activities that use resources. To the degree that the same actors perform multiple coordinating activities or produce resources employed by coordinating activities, they can be identified as important coordinators within a process. MAG account executives produce both quotes and instructions (resources) while performing 5 of the nine coordinating activities identified in the prior step. From this perspective they appear to play an important organizational role in supporting coordination.

Dependency aggregation. The three steps of activity-focused analysis have so far suggested that (1) preparing quotes and instructions are at least coordination-intensive activities; (2) producing personalized mailing materials is probably a critical process step around which coordination activities cluster; and (3) AEs perform much of the coordination required to define and complete profitable jobs. These suggestions focus attention on the potential coordination provided by processes related to preparing quotes and setting up jobs.

If these coordination activities manage dependencies, it is reasonable to move one level higher in the process hierarchy and ask what coordination they perform. Doing so considers the relationship between the company and its customers as the company runs a mailing job (e.g., "Run Job"). From this perspective, it appears that "Prepare Quote" manages the usability of a job to the customer and the profitability of the job to MAG: in other words, it manages usability constraints. "Set Up Job" appears to manage the sequencing of activities within MAG operations (recall that three sets of instructions are prepared and distributed, one to each functional area, that direct how the functional areas are to interact during the job). In this sense, "Set Up Job" is managing prerequisites.

To summarize, in activity-focused analysis, we first search for activities that appear to be examples of coordination mechanisms, and then check for dependencies that are managed by these activities. This bottom-up approach offers an alternative view of the process that is complementary to the results produced by a top-down, dependency-focused perspective. Where the results overlap the two analyses offer the means for producing confirmatory evidence of coordination choices made by the organization. Either approach can confirm or disconfirm process characteristics suggested by the other. Dependency-focused analysis proceeds from the perspective of high-level goal structure, while activity-focused analysis begins with paper flows and process artifacts that exist deep within the organization.

Step 6: Verifying a model

Process models may be as valuable for the insights that are developed in the process of building them as for the final process diagram. To this end verification plays a particularly important role in the techniques suggested here. We suggest two verification techniques in particular as complementary mechanisms for improving process representation and analytical accuracy.

The first of these is the negative case method [11]. Candidate process representations are developed and discussed to discover what is missing in the representation of the process. Gaps and ambiguities identified guide further data collection. These omissions can be identified by discussing the process model as it evolves with the actors who are involved in the process—discussions which often trigger the need to go back and revisit various process representation decisions made earlier.

Triangulation provides a second opportunity to verify the faithfulness of process representations. We use the term to refer to the ways in which process models are discussed with process actors. The emphasis here complements the negative case method by assessing what process elements appear to generate broad representational agreement from groups of process participants. Representational accuracy, in this sense, can be corroborated by broad agreement among process participants.

Negative case analysis and triangulation are included here to highlight the importance of internally consistent verification of process models, context descriptions, and analyses by the individuals who participate in the process themselves. Since any organizational process is open to differing interpretations by each of its participants [8], and since even the problems that processes are designed to "solve" are open to definition-by-interpretation [20], subjective verification may be the best consistency-control available to field teams.

From a research design point of view, the dangers of subjective verification, even by multiple respondents, are well known [21]. It is important to recognize, however, that managers and other professionals working within organizations face the same limitations in understanding observed behavior: for them, low-level subjective consensus represents one key mechanism whereby groups jointly interpret events. Iterative, multi-source verification (e.g., of process descriptions) in this view represents the same level of reliability that process designers themselves must handle in actual practice. To the degree that an expanding understanding of process characteristics can add structure and consistency to process representations, the reliability of process descriptions can be improved. Absent this, however, the use of iterative, multi-source verification seems not only an achievable means for checking descriptive fidelity but also one that very appropriately reflects real-world conditions.

Summary

The result of the six step process outlined above is documentation of a process that includes activities, actors, resources and dependencies between them, as well as identification of how the dependencies are currently managed. This process documentation can then be subject to test to ensure it is reasonable, that it makes sense to people or can be used to communicate the process. In the next section we will discuss how such documentation might be useful as a basis for process improvement.

Using dependency analyses as a basis for process improvement

Our main purpose in creating process representations is to support process improvements. Documenting the dependencies and coordination mechanisms of a process provides an approach to developing new processes. New activities can be proposed to manage poorly managed dependencies or alternative coordination mechanism can be considered to manage each dependency. Note that mechanisms are themselves activities, with their own set of dependencies. Replacing one mechanism may therefore eliminate some problems while creating an entirely new set to be managed.

When the custom mailing process began to break down, MAG managers responded in ways suggested by their experience with traditional, high-volume, standardized mailings. They initially fixed their attention on lowering costs by trying to make custom jobs run as smoothly as the standardized non-custom work. They focused on rearranging the company's internal production processes (e.g., the steps by which MAG sorted data and prepared mailings). Unfortunately, this perspective meant that the cures initially suggested for the custom business proved worse than the disease, as it missed the need for flexibility demanded by customized contracts.

Our analysis suggest ways in which to modify coordination strategies and the tools used to implement those strategies. For example, dependency-oriented analysis might to suggest ways in which

AEs can redesign the quoting process to reduce project lead-times. Activity-based analysis, however, can contribute useful detail describing how to redesign a quotation *form* to be used in describing the cost of services to clients over the telephone. Taken together, the two approaches can contribute guidance to information systems development designed to resolve the timing and flexibility problems identified as threats to the custom business. Dependency-focused analysis can contribute to clarity of process purpose, while activity-focused analysis can contribute insight about implementation detail.

The process analysis performed for MAG services as part of this project led to prototypes of process improvements. After completing the analyses described above, the organization developed software based on a commercially available groupware package to experiment with making three changes in the Run Job process.

First, the software provided AEs with cost estimates for any combination of mailing tasks, using an interface that enabled them to build accurately costed job quotations during a telephone conversation. The same software generate quotation letters semi-automatically. Second, the system fed data electronically to instruction forms. These forms were extremely simple in design, and accumulated operational detail only for the tasks specifically required in any one project. Lastly, the system provided a series of checkpoints so that AEs could electronically monitor job progress, enabling them to report back to customers on a daily basis if necessary. All three of these innovations provided a better way to ensure that the work done is what the customer requested and is correct, that is, to manage the usability portion of the high-level flow dependency discussed above. AEs reported that these design changes had the potential for increasing their capacity for custom work. At the end of our study, the company was considering whether to develop a commercial version of the system.

A coordination perspective also provides some insight for goal resolution. To the degree that usability constraints threaten to be incompatible, the process coordinating them may include explicit activities devoted to resolving potential conflicts. In many organizations, for example, order-taking begins to resemble sales negotiation as activities are added to ensure that the order will be usable for both producer and consumer. In engineering-intensive businesses, the RFP (request for proposal) process can be understood in this fashion. In simpler businesses, as well, an important coordination opportunity often arises as new business enters the workflow. At MAG Services, for example, a quoting process controls how orders for customized services are placed.

More specifically, the perspective provides a mechanism for summarizing potentially divergent goals surfaced by MAG's business relationships. For example, it suggests how a customer's purpose might interact with the organization's internal goals. A customer is likely to seek to generate the maximum number of useful sales leads in the shortest available time. MAG has a need for profitable growth. In coordination terms, these goals represent *usability* constraints affecting the flow that MAG coordinates. Each mailing service, this representation implies, must remain usable to both the customer and to the supplier—i.e., it must provide sales leads that satisfy some range of customer criteria yet remain profitable to MAG.

Tradeoff matrices

As well, the models can serve as a basis for articulating the tradeoffs available between different versions of a given process. Tradeoff matrices contrast process characteristics across different versions of a process. Consider the variations of the "Run Job" activity illustrated in Table 3. One version of the process refers to selling traditional business; the second refers to selling custom jobs. The dependencies underlying the coordination analyses above apply to both but the strategies employed for coordinating these dependencies differ. The tradeoff matrix in Table 4 suggests some of the ways in which the two types of business compare. The comparison suggests ways in which process performance can be improved by redesigning the ways in which process dependencies are coordinated.

For example, the trade-off matrix suggests that coordinating custom projects is more time-sensitive than initiating non-custom work. This comparison implies that if the company could facilitate AEs' and clients' understanding of the cost implications of custom services, the logistics and productivity of MAG's custom services might be enhanced. Analyzing usability, transfer, and prerequisite dependencies suggests a range of alternatives for improving communication between AEs, customers, and operations staff in all stages of job definition and execution, extending from technology-intensive solutions such as an on-line job definition system that allows new customers to design their own customized service to relationship-intensive solutions such as pairing operating staff with AEs in custom-project teams.

Discussion

To put our contribution into perspective, we will conclude by briefly comparing our work to other process analysis techniques and evaluating our technique.

Comparison to other process analysis techniques

Process design and coordination problems have been approached from diverse perspectives, including economics, organization theory, computer science, ecology, and general management theory. We will briefly review alternative approaches.

Perhaps the simplest form of a process description is a simple verbal account. Such accounts are commonly used and have the advantage that little or no special training is needed to produce or understand them. However, there are two key problems: first, it is difficult to check a verbal description for completeness or consistency; second, verbal descriptions do not easily suggest the space of possible improvements. Therefore, most analysts use a more formal representation, as do we. It is interesting to note that soft systems methodology uses “all the verbs in the English language” [1, p. 164] for building conceptual models, but the goal of these models is to “generate radical thought” (p. 170) and deliberately not to be descriptions of the actual system.

A PERT chart provides a detailed representation of a process, specifying the exact activities taken, when they begin and end, sequence dependencies between activities and even which actors or resources are involved with which activities. PERT charts have one major drawback for the purpose of process improvement: they usually are used to present or plan a single execution of a process and do not represent the range of possible alternatives. However, our representation captures many of these details, such as dependencies between activities and the use of resources.

Managers and analysts interested in improving processes often use some version of a flowchart to represent process characteristics. Flowcharts drop some information in a PERT chart, but still indicate the activities to be performed, the order in which they are performed, and may include information on who does each activity or how long an activity takes to perform. However, they are not especially good at suggesting alternative activities that accomplish the same ends, at demonstrating feasible alternative activity sequencing, or at projecting what changes might be required if different actors performed selected activities.

Our representation is most similar to a dataflow diagram, which represents the steps of a process but focuses on the ordering relationships imposed by the fact that data produced by some steps must be used by others [e.g., 22]. Many data-flow techniques, such as IDEF0 and SADT, include decomposition as a key aspect. These representations are similar except they do not represent the full range of dependencies nor explicitly note the coordination mechanisms.

To represent processes involving multiple actors, we may want to focus on the interactions between the actors. One approach to modelling interacting processes is suggested by *Petri nets* [14] and various representations derived from them [e.g., 10; 17]. A Petri net is similar to a finite state machine, but allows multiple states to be "marked" simultaneously. Transitions between states may be synchronized, since multiple states may have to be marked at the same time for a particular transition to occur. To the extent that the activities we model have multiple inputs, then our representation can be seen as equivalent to simple Petri nets, although we do not take advantage of that fact.

A second approach to representing multiple actors is to represent the process followed by each individual separately, using any of the techniques described above, and explicitly modelling the exchange of information or objects between them. For example, the modelling technique developed by Crowston [3] represents individual actors as programs written in logic. These actors can perform a variety of actions to achieve their goals, including speech actions to change the states of other actors. We believe that such representations could be used as a basis for simulating processes, thus providing a more detailed approach to examining tradeoffs.

Evaluation

While the technique proposed in this paper embodies a theory, it does not provide a way to test the theory. Indeed, Checkland [1] argues that methodologies can not be proven to work (or not to work) in the scientific sense (p. 241). Instead, the evaluation of the technique rests on how well it accomplishes the two goals we set out in the introduction: generativity and ease of use. The technique is a success if analysts can use it and if they find it provides insights into the process and how to change it. These two tests might partially trade-off against each other: for example, if the technique provides unique insights, then analysts might be willing to undergo more of a learning process.

We do not claim to have done quite as thorough a test of our technique as Checkland [1, p. 193] who offers the results of more than 100 attempts to use soft systems methodology. However, in the past three years, one of the authors has taught coordination theory and the methodology to four courses and two project teams, totalling approximately 70 masters-level students. Learning and applying our methodology occupied six to eight weeks of each course. Approximately half of these students were managers employed in a variety of industries and enrolled in a part-time MBA program. Students were required to redesign a process based on their analysis and to develop an information systems prototype that would support that redesign. Over all, they completed more than 40 process design projects, based on observations and analyses at large and small companies. The result was fairly consistent process innovation that exceeded the expectations of project participants. Clearly these experiences, while suggestive, do not prove much about our technique. We are therefore beginning a study where the technique will be tested more directly.

Our teaching experiences underline the paradigm shift need to think about organizational processes. It has been difficult for participants to make the transition from focusing on inputs (e.g., strategies and resources) and outputs (e.g., organizational results) to focusing on the processes that derive those outputs. As well, dependency analysis has been the most confusing aspect of methodology, hence our focus on it in this paper.

Stated alternately, coordination theory does not make strong predictions about what should happen to any single organization. Rather than the specific accuracy of its predictions, therefore, an appropriate test for the theory is its utility for organization designers. Our approach is a success if those attempting to understand or redesign a process find it useful to consider how various dependencies are managed and the implications of alternative mechanisms.

Conclusion

To conclude, we will briefly discuss the implications of our focus on dependencies for the design of analysis tools and the practice of managers and other process analysts.

Suggestions for design of tools

The approach has strong implications for the design of process analysis tools. Many CASE tools can represent a decomposition hierarchy of activities. Some could handle the dependencies linking activities and resources. However, none that we know of explicitly represent the link between coordination mechanisms and the dependencies they manage. For example, to assist analysts building such representations, it would be handy to be able to drag an activity on to a dependency to indicate that the dependency is managed by that activity or to click on a dependency and pop up the list of alternative specializations.

Since dependencies arise from shared use of resources, the representation of an activity could include an indication of the resources they need from which the system could automatically figure out some dependencies. For example, if two activities need the a resource of which there is only one known instance, then the system might suggest resource sharing mechanisms; if there is no known resource, it can suggest resource procurement mechanisms. If the resource is an actor, then a task assignment mechanism is needed. As an example, we are currently using these techniques to compile a handbook of organizational processes at a variety of levels and in different domains [13]. Managers or consultants interested in redesigning a process could consult the handbook to identify likely alternatives and to investigate the advantages or disadvantages of each. Coordination theory makes the handbook feasible by providing a framework for describing more precisely how processes are similar and where they differ.

Implications for practitioners

Even though many people have documented and studied organizational processes, our approach to this problem is novel in important ways. Most importantly, in analyzing a given process, we identify the key activities that must be performed for the goal to be achieved, the resources created and consumed by these activities, and the dependencies between them. We define the managing of these dependencies as the coordination activities, and we postulate that there will be a set of generic coordination processes (and their various specializations) that will appear over and over in different processes.

By identifying the various types of dependencies and the generic processes for managing them, we believe that we can create more concise process descriptions. A second, and potentially more important benefit, however, is that this approach can help us generate new possibilities for processes. If we know that, in general, there are several possible coordination processes for managing a given dependency, then we can automatically generate all of them as possibilities for managing that dependency in any new process we analyze. Some of these possibilities may be new or non-obvious, and their generation requires no specific knowledge of the process other than the type of dependencies it involves.

The choice of coordination mechanisms to manage these dependencies results in a variety of possible organizational forms, some already known and some novel. The relative desirability of mechanisms is likely to be affected by the use of new information systems. For example, the use of a computer system may make it easier to find existing solutions to a problem, either in a database or from geographically distributed coworkers. Such a system could reduce both duplicate effort and coordination costs.

References

[1] P. Checkland. (1981). *Systems Thinking, Systems Practice*. New York: Wiley.

[2] K. Crowston. (1997). A coordination theory approach to organizational process design. *Organization Science*, *8*(2), 157–175.

[3] K. Crowston. (1991). Modelling coordination in organizations. In M. Masuch and G. Massimo (Eds.), *Artificial Intelligence in Organization and Management Theory*. Amsterdam: Elsevier.

[4] K. Crowston. (1996). *A taxonomy of organizational dependencies and coordination mechanisms* (Unpublished manuscript). Syracuse University School of Information Studies.

[5] T. H. Davenport. (1993). *Process Innovation: Reengineering Work Through Information Technology*. Boston: Harvard Business School.

[6] T. H. Davenport and J. E. Short. (1990). The new industrial engineering: Information technology and business process redesign. *Sloan Management Review*, *31*(4), 11–27.

[7] B. G. Glaser and A. L. Strauss. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago: Aldine Publishing.

[8] J. R. Hackman. (1969). Toward understanding the role of tasks in behavioural research. *Acta Psychologica*, *31*, 97–128.

[9] H. J. Harrington. (1991). *Business Process Improvement: The Breakthrough Strategy for Total Quality, Productivity, and Competitiveness*. New York: McGraw-Hill.

[10] A. W. Holt. (1988). Diplans: A new language for the study and implementation of coordination. *ACM Transactions on Office Information Systems*, *6*(2), 109–125.

[11] L. H. Kidder. (1981). *Research Methods in Social Relations* (4th ed.). New York: Holt, Rinehart and Winston.

[12] T. W. Malone and K. Crowston. (1994). The interdisciplinary study of coordination. *Computing Surveys*, *26*(1), 87–119.

[13] T. W. Malone, K. Crowston, J. Lee, B. Pentland, C. Dellarocas, G. Wyner, J. Quimby, C. Osborne, A. Bernstein, G. Herman, M. Klein and E. O'Donnell. (in press). Tools for inventing organizations: Toward a handbook of organizational processes. *Management Science*.

[14] J. L. Peterson. (1977). Petri nets. *ACM Computing Surveys*, *9*(3), 223–252.

[15] J. Pfeffer and G. R. Salancik. (1978). *The External Control of Organizations: A Resource Dependency Perspective*. New York: Harper & Row.

[16] E. H. Schein. (1987). *The Clinical Perspective in Fieldwork*. Beverly Hills: Sage.

[17] B. Singh and G. L. Rein. (1992). *Role Interaction Nets (RINs): A Process Definition Formalism* (Technical Report CT–083–92). MCC.

[18] G. F. Smith. (1988). Towards a heuristic theory of problem structuring. *Management Science*, *34*(12), 1489–1506.

[19] J. P. Spradley. (1979). *The Ethnographic Interview*. New York: Holt, Rinehart, Winston.

[20] K. E. Weick. (1969). *The Social Psychology of Organizing*. Reading, MA: Addison-Wesley.

[21] R. K. Yin. (1984). *Case study research: Design and methods*. Beverly Hills, CA: Sage.

[22] E. Yourdon. (1989). *Modern Structured Analysis*. Englewood Cliffs, NJ: Yourdon.

Figures and tables

Table 1: A preliminary typologies of dependencies and coordination mechanisms

Task uses resource

- Decomposition 1
  - determine needs = manage usability from acquiring resource to using
  - identify resources
    - ads
    - prepared list
    - only one resource?
  - collect information on resources
    - by bidding
    - manager knows
  - pick best
  - do assignment
    - mark resource in use
  - manage flow dependencies from acquiring resource to using resource

- Decomposition 2
  - determine needs
  - plan tasks to accomplish needs
    - decompose the goal into subgoals
    - pick a set of activities that achieve the goals
    - integrate if necessary

Multiple tasks use common resource

- ensure same version of sharable resources
  - destroy obsolete versions
  - copy master prior to use
  - check versions prior to use
  - detect and fix problems after the fact
- schedule use of non-shareable but reusable resources
  - check for conflict before using and then mark the resource as in-use
  - manage flow of resource from one task to another
- allocate non-reusable resources
  - divide the resource among the tasks
  - abandon one task
  - get more resources

Table 1 continued.

One task uses a resource created by another

- Decomposition
    - usability (i.e., the right thing)
        - user adapts to resource as created
        - creator gets information from user to tailor resource
        - 3rd party sets standard followed by both
        - (NB. Fit = multiple usability dependencies)
    - prerequisite (i.e., at the right time)
        - producer produces first
            - follow plan
            - monitor usage
            - wait to be asked
                - standard reorder points
                - when out
                - just-in-time
        - consumer waits until produced
            - monitor
            - be notified
    - accessibility (i.e., in the right place)
        - physical goods
            - truck
        - information
            - on paper
            - verbally
            - by computer

Multiple tasks create the same resource

- Decomposition
    - Detect common output
        - database of known problems
    - Manage
        - overlap or same
            - eliminate one task (manage shared resource)
            - merge tasks take advantage of synergy
        - adds are incompatible
            - abandon one
            - add a new goal to fix it up
            - don’t try to achieve them at the same time

Table 2. Actors in the MAG case.

Class	Actor	Description
Customer	Customer representative	The contacts at the client company who represent the client contacts; usually staff members in a Marketing department.
Sales	MAG Salesperson	Salesperson who initiates contact with new customers. Works in a different department than the AEs.
Operations	MAG Account Executive (AE)	Account Executive who quotes, schedules, and manages traditional and customized mailing jobs.. Also initiates new jobs from existing customers.
Data Entry	Data Entry Supervisor	Schedules daily data entry work.
	Data Entry Staff	Keypunch incoming inquiries.
Technical Services	Database production mgr.	Executes data transformations to be completed on job-specific data using MAG-proprietary database.
	Programmer	Provides job-specific programming for custom jobs.
Mailing Services	Batch dispatcher	Transfers jobs to packing/sorting tables.
	Packer	Pick, pack, and sort mailings.

Table 3. Resources used in the MAG case.

Resource	Linking Actor(s)	Description
Job Quote	Customer, AE	Defines scope of work approved by customer.
Data Entry instructions	AE, Data Entry	Instructs staff on data entry details.
Batch(es)	Data Entry, Technical Services, Mailing Services	Organizes the inquiry stream into batches that are processed by MAG's database and packed/shipped as separate units.
Technical Services Work Order	AE, Technical Services	Instructs database production managers on how to sort and/or modify data.
Mailing Services Work Order	AE, Mailing Services	Instructs mailing services staff how to pick, pack, and sort job.
Mailing shipment	Technical Services, Mailing Services	Physical components of mailing shipment as they are produced off bulk printers, picked, packed, and shipped.

Table 4. Comparison of custom and non-custom work.

Service	Volume	Duration	Processing	Task Order	Deadlines	Contact
Custom	Potential	3-4 months	Custom	May vary	Tight	Daily
Non-custom	Certain	Year(s)	Standard	Standard	Set by MAG	Quarterly

Table 5. Summary of initial analysis.

Question	**Answer**
Process boundaries	Provide MAG Service and its decompositions.
Process goal	To generate revenue by selling company services that satisfy the needs of mailing services customers. In this sense the goal can be interpreted as providing services that convert a customer with a marketing need into a client who with a successful mailing-based marketing campaign.
Process outputs	Direct mailing services on behalf of client.
Process inputs	Client need for mailing services (as defined by client approval of a specific set of mailing services).
Resources	Salespeople, account executives, production staff, mailing materials, warehouse space, computer equipment, printing equipment, temporary staff, work orders, and instructions.

Table 6. Dependency-focused analysis—coordination activities.

Activity within "Run Job"	Description	Purpose	Constraint managed
Prepare quote	The AE ensures that MAG is producing the exact tasks that will deliver what the customer wants.	Produce the right service	Usability
Set up job	The AE prepares and distributes instructions that describe, in detail, what each functional unit with MAG must accomplish to complete the job.	Produce the right service at the right time.	Prerequisite
Produce mailing service	Operations completes data entry, produces the mailing, ships required materials, and distributes market data back to the client.	Move the required product to the right place.	Transfer

Table 7. Dependency-focused analysis—coordination strategies.

Dependency	Between	Managed by	Key attributes	Coordination strategy	Implications for custom jobs
Flow	Use Mailing Service, Use Market Data and Sales Leads	Provide MAG Service	Goals: Satisfied client, profitable job. Input: Job quote Output: Mailing service Resources: AEs, staff, database system, mailing system	Run Job	Jobs coordinated using strategies developed for traditional, non-custom work are breaking down, leading to missed deadlines, below-target profitability, and lost business.
Usability		Prepare Quote, Provide Status Reports	Key variables: Service configuration Lead times Costs	Negotiate quote with customer	Custom jobs are often so complex that the customer does not know what to ask for.
Prerequisite		Set Up Job	Key variables: Service specification Set up time Accuracy	Prepare and distribute instructions to all operational departments	Instructions for traditional work are one-way. For custom jobs, AE does not get feedback about mistakes until it is too late.
Transfer		Produce Mailing Service	Key variables: Delivery speed Quality Accuracy	Process leads, pick, pack and ship mailing materials	Reporting systems were designed on a batch basis for standard work; AE cannot unravel mistakes without long delays.

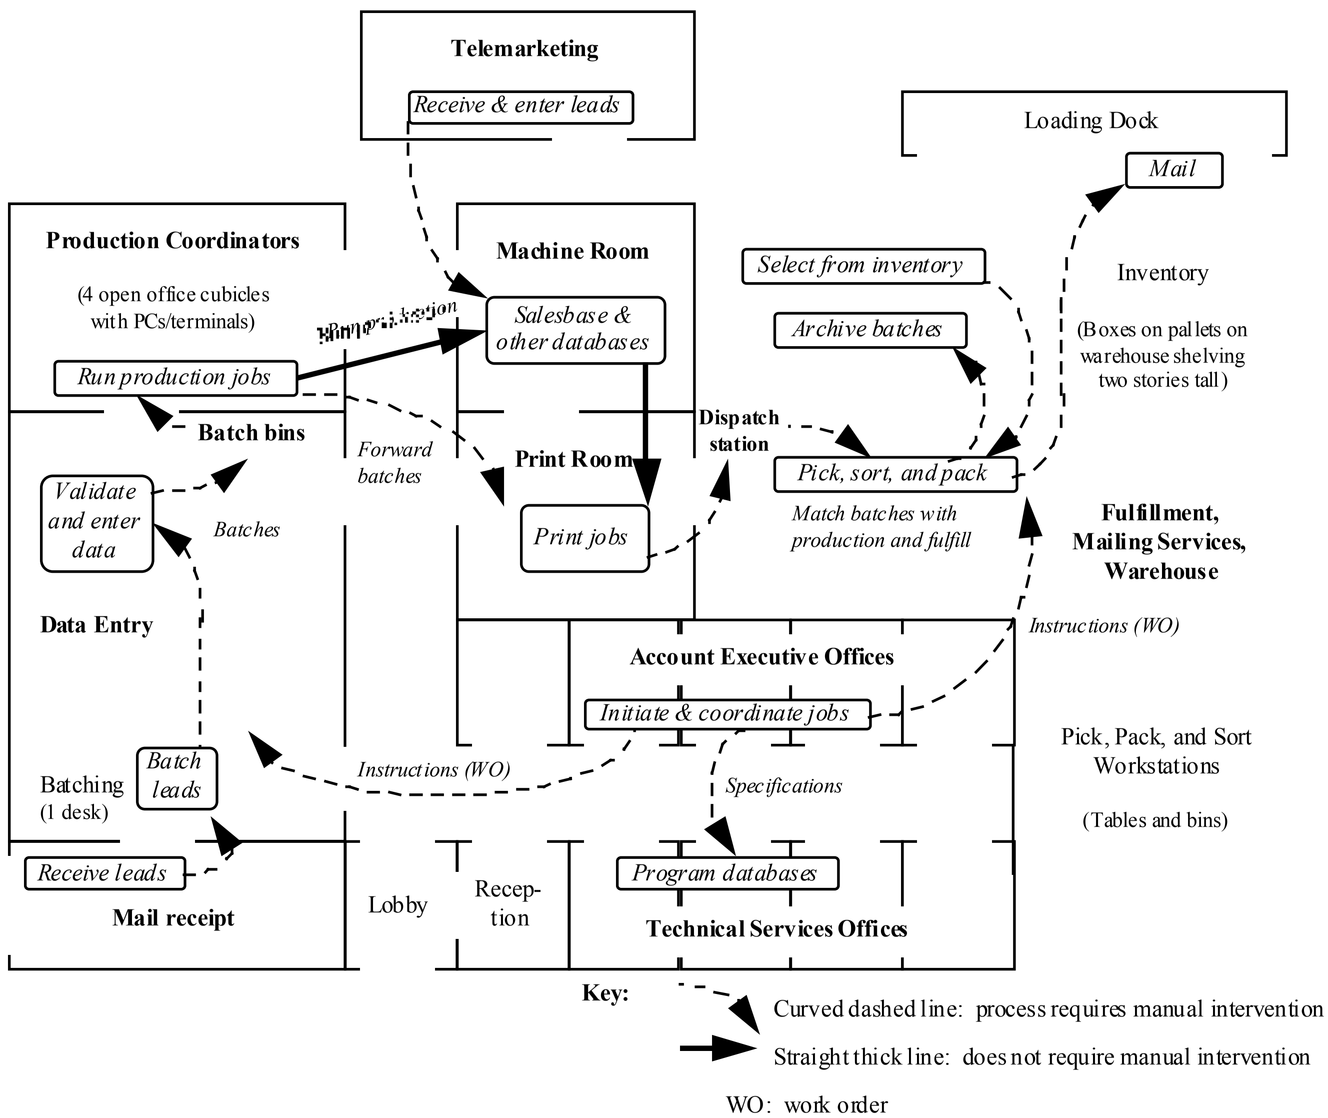
Telemarketing
Receive & enter leads
Loading Dock
Mail
Production Coordinators
(4 open office cubicles with PCs/terminals)
Run production jobs
Machine Room
Salesbase & other databases
Select from inventory
Archive batches
Inventory
(Boxes on pallets on warehouse shelving two stories tall)
Batch bins
Forward batches
Print Room
Dispatch station
Pick, sort, and pack
Match batches with production and fulfill
Validate and enter data
Batches
Print jobs
Fulfillment, Mailing Services, Warehouse
Data Entry
Instructions (WO)
Account Executive Offices
Initiate & coordinate jobs
Pick, Pack, and Sort Workstations
Batching (1 desk)
Batch leads
Instructions (WO)
Specifications
(Tables and bins)
Receive leads
Mail receipt
Lobby
Recep-tion
Program databases
Technical Services Offices
Key:
Curved dashed line: process requires manual intervention
Straight thick line: does not require manual intervention
WO: work order

Figure 1. Basic work flow at MAG services.

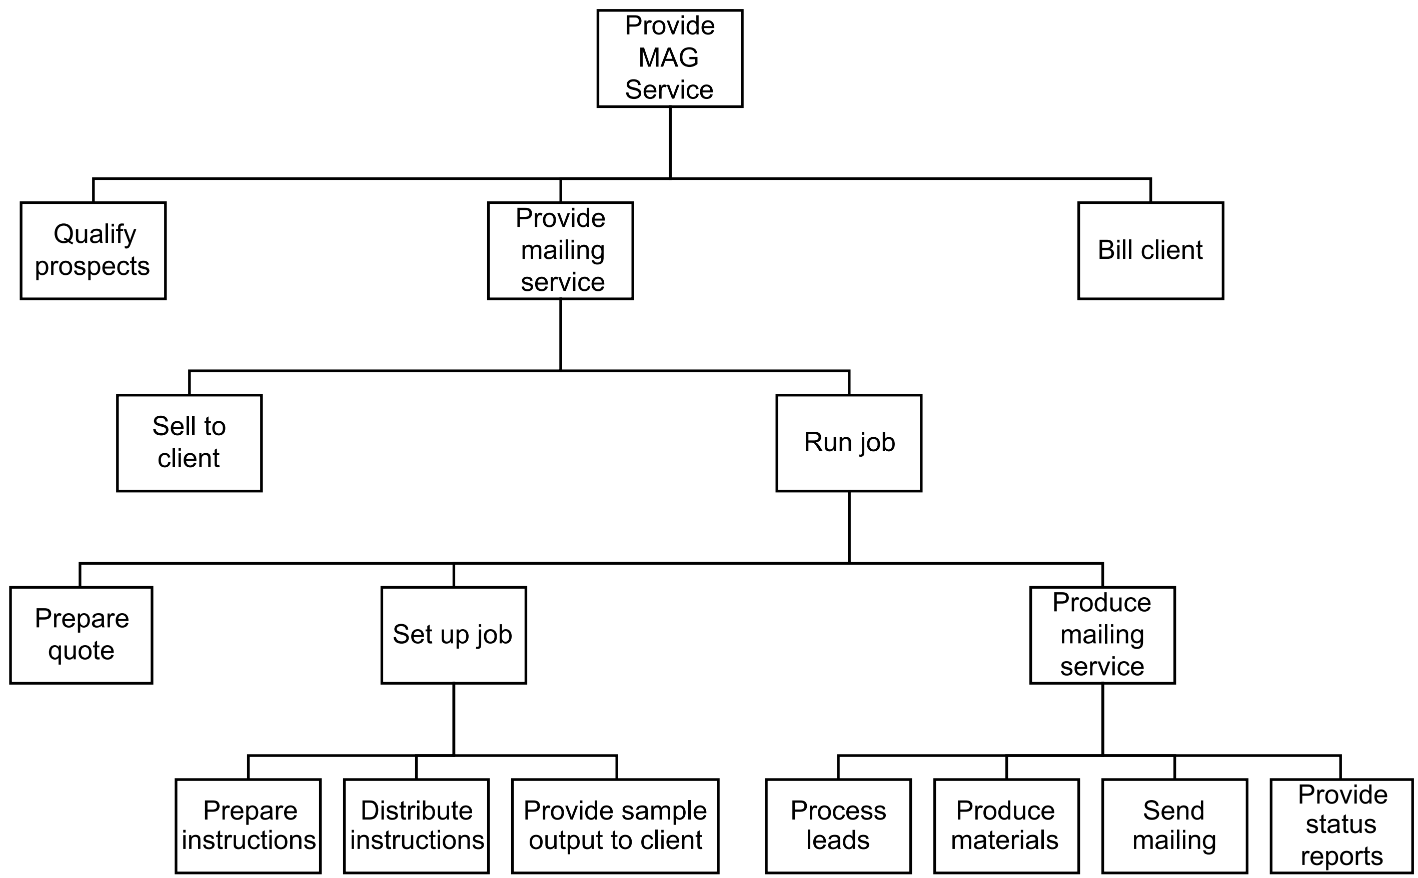

Figure 2. High-level process decomposition view.

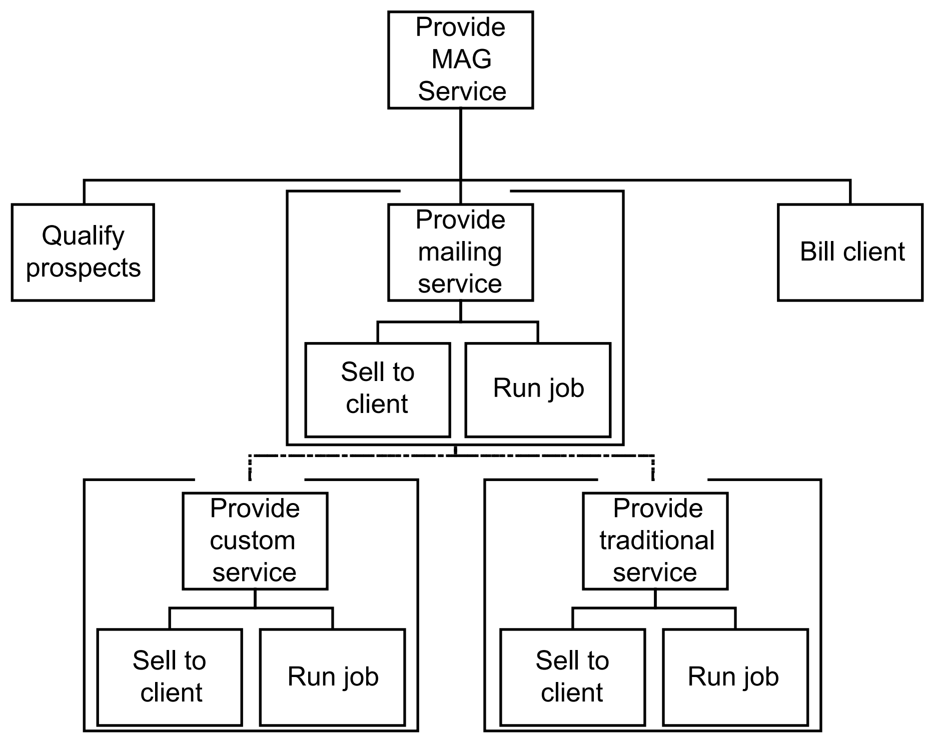

Figure 3. Specializations illustrate process variety.

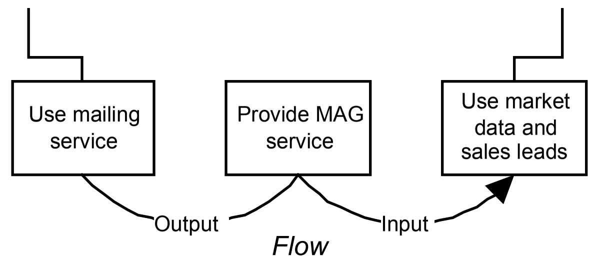

Figure 4. MAG Services as a step in a value chain.

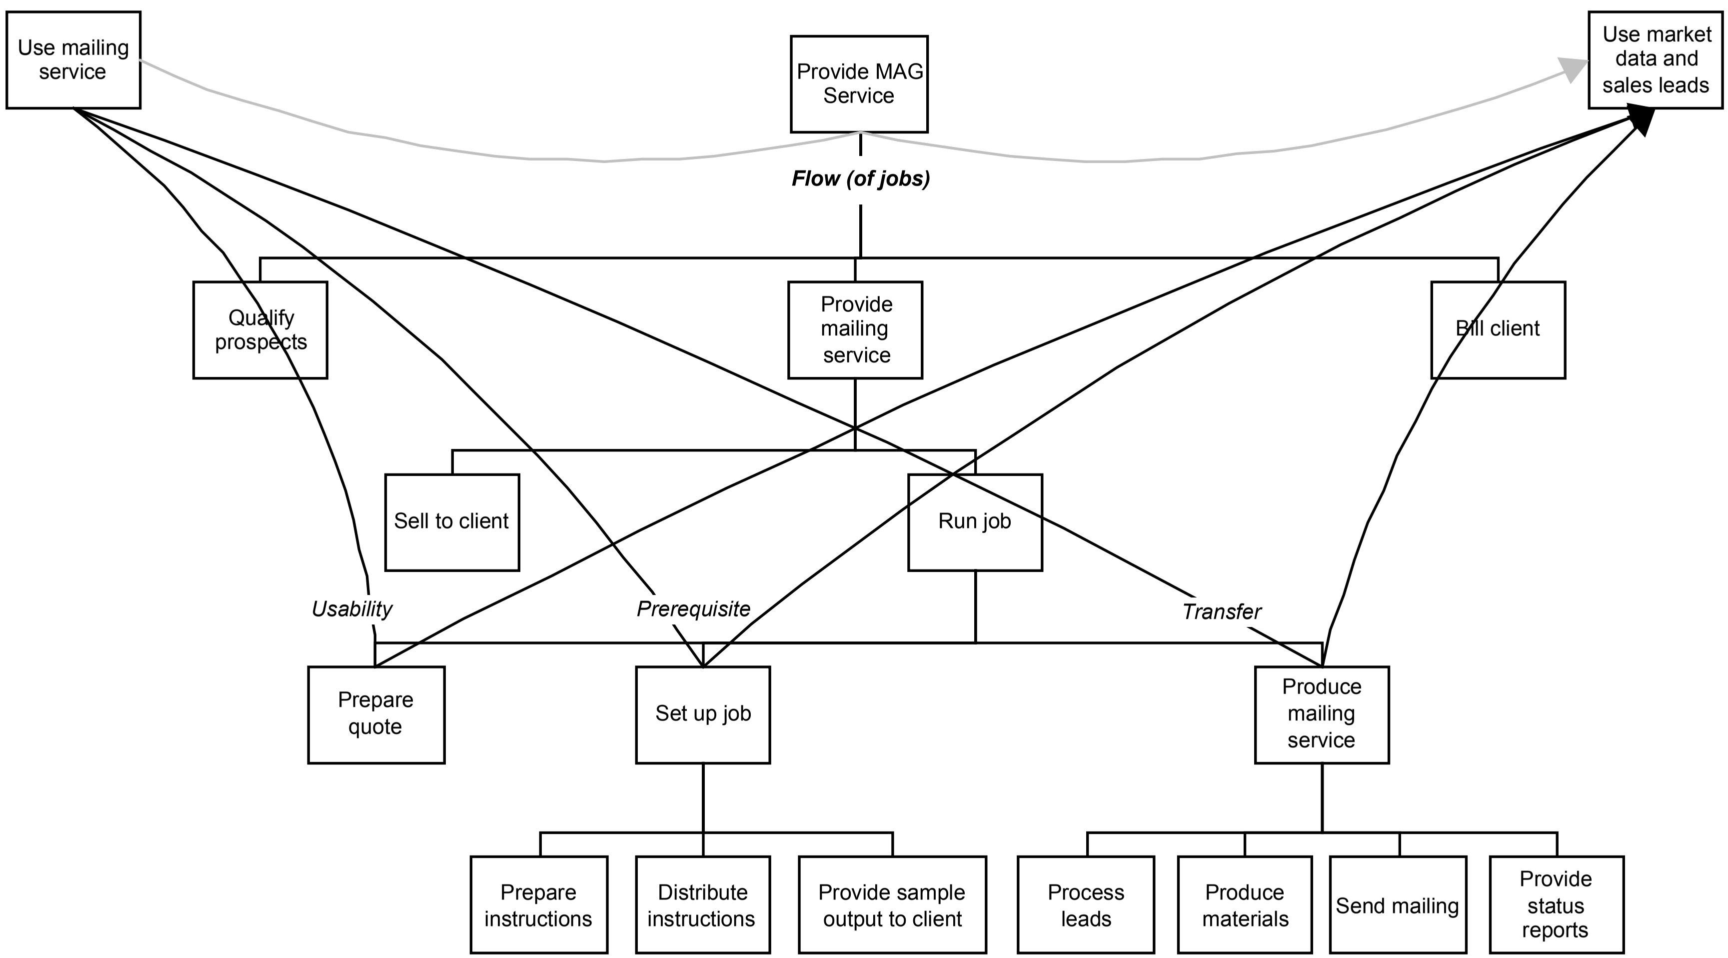

Figure 5. Coordinating subdependencies within the Run Job process.

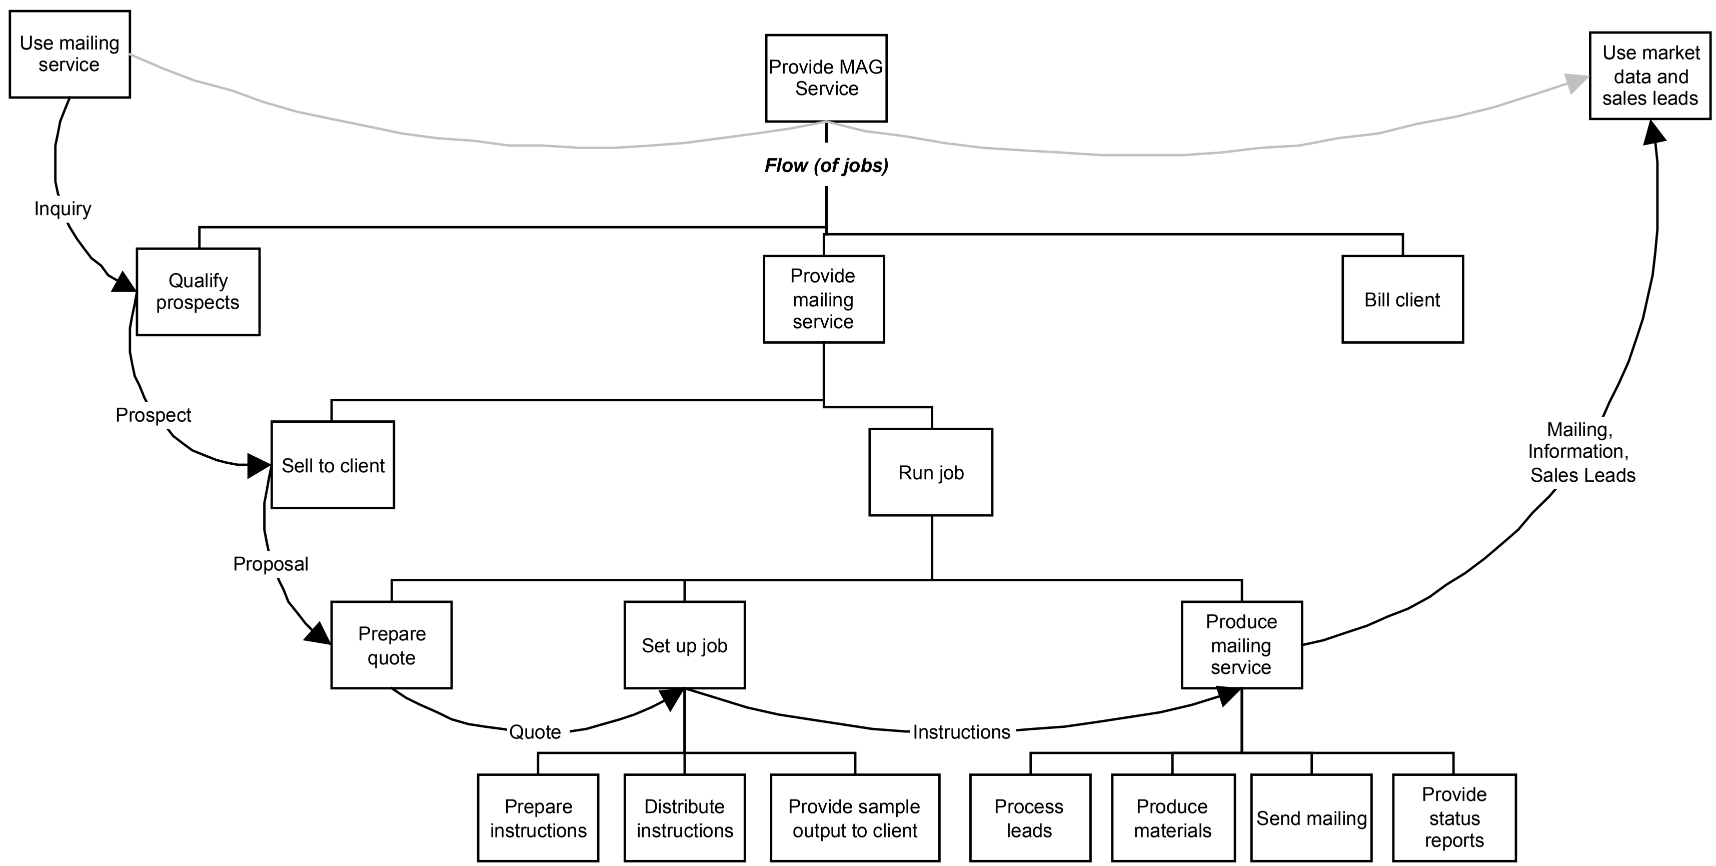

Figure 6. Paper flow and resources at MAG Services.