"Yeah, the Rush ain't here yet -- Take a break": Creation and use of an artifact as organizational memory

Christine A. Halverson
IBM T.J. Watson Research
krys@us.ibm.com

Mark S. Ackerman
University of Michigan
ackerm@cs.umich.edu

Abstract

In order to understand organizational memory, it is important to understand how things become adopted as memory resources in organizations. In this paper, we describe the genesis and use of an artifact that became a memory resource for a wide range of activities. We discuss how the creation and use of the rush cheat sheet (RCS) and its associated representations at Dallas Ft. Worth TRACON brought together information and expert knowledge across organizational boundaries. Multi-organizational information became synthesized in a composite that could be used as a resource by the contributing organizations, acting as a boundary object. However, it is multiple representations of the same data that enable it to be so used. Using distributed cognition theory, we examined the conditions under which data transforms from an internal resource to a boundary object; speculating about domain generalization.

1. Introduction

Information reuse is clearly possible within organizations – every organization does it every day. Yet, as much as it occurs, there remains a burning sense that more and more information is slipping through our collective fingers. A naïve sense of reuse starts with the belief that there must be information available to be reused and one should know how to do so. This belief has found a new home in terms of the objectives for organizational data stores. Yet, the promise seems much larger than the reality.

Accordingly, we would like to know how reuse is accomplished as an everyday activity [18]. To do this, we would like to understand how organizational memory comes into service for an organization. Indeed, various artifacts – computerized or not – have always been used to help record, and remember the details of work activity. However, surprisingly little is known about the details of such reuse within organizations. Cognitive scientists have looked at artifacts as a way to aid individual memory (e.g. [20,27]). Similarly, anthropologists, sociologists, and historians have looked at social structure as a way to encode process (e.g., [29]). Our question is how some artifacts can act to encode both memory (or content) and process in an organization.

This paper examines organizational information reuse by tracing the life cycle of a relatively low-tech artifact used in air traffic control (ATC). This artifact – called the rush cheat sheet (RCS)—began as the product of a search for information. Over time it has become, in its own right, a source of expertise and information—not only for its original purpose but also for a range of originally unintended inquiries. While the details of this analysis are clearly limited to the domain, their overall lessons have broader applicability.

We begin by situating our approach with respect to other investigations. We lay out the analytical background that motivates tracing this artifact, including the analytical approach and constructs applied from previous work. We then introduce the domain and describe how the RCS came to be and how its use has changed over time. We discuss several aspects of the RCS that affect its use, and that further distinguish how artifacts come to be used as organizational memory. For example, the mixed provenance of its origins may contribute to its successful use on both sides of an organizational boundary, and follow that with a discussion of what this might mean about the conditions under which similar artifacts would be successful.

2. Related work

Much of organizational memory literature has focused on one of three aspects. First, there is broad theorizing about organizational memory in terms of general frameworks (e.g., [10,16,33,34,37]). A second aspect covers the design of organizational memory systems, such as Answer Garden [2-4, 8] . While many of the system building papers include studies of systems under use (e.g., [11, 13,21,23,26,39]), what is still lacking is the detailed study of organizational memory in the wild. Recognition of this lack has now resulted in a third area – finely detailed analysis of field studies of organizational memory (e.g., [17,28,30,31]).

Our focus here is in the last area. The design of Answer Garden and studies of its use led us to undertake more detailed studies of organizational memory in a hotline setting [5-7]. That work established a general analytical baseline for examining organizational memory and information reuse, using distributed cognition theory [19]. These analyses proposed critical analytical concepts that we believe are necessary for understanding information reuse. These are *boundary objects, mixed provenance,* and *trajectories of use*. In the section that follows we introduce distributed cognition theory and discuss these constructs.

3. Our Approach

3.1. Distributed Cognition

Our analyses are based on the work of Hutchins and colleagues [14,15,19]. Hutchins' *distributed cognition theory* is focused on how cognitive systems are organized and operate, thus re-situating cognition in its socio-cultural context. In this sense it is part of the recent recognition that societies and organizations demonstrate cognitive properties that are different from those of individuals. [12,19,32]

Distributed cognition theory serves as a useful conceptual framework for the analysis of human organizational systems, and gives us a way to approach the processes concerned with information use and reuse. Looking at the specifics of how information is represented, and where it is represented, highlights how both storing and retrieving information can be affected. Generalizing across domain examples [14,15,19,20] there are three dimensions across which cognition tends to be distributed: people, time, and representations. All three can impact information reuse.

Cognition is distributed through *time* in such a way that the products of earlier events transform the nature of later events. Culturally defined routines, such as the "right way" to do something, are examples of this. This temporal distribution also serves to distribute the cognition across *people*. A craft master may have devised a new way to make dye, but as the recipe gets passed along, generations of assistants need only know how to mix the components accurately.

Similarly, artifacts can, and often do, embed a history of their use in the form of their *representations*. A ruler is marked with lines having meaning that affects their current use. Saying that a piece of paper is 8 ½ by 11 inches assumes that the measurement, as well as the number system, is meaningful to both the speaker and the hearer. That meaning is based on their previous use and the history of how those markers came to have meaning with respect to socially agreed measurements. Furthermore, in the US that is interpreted as standard "letter size", while in Europe it is an oddity compared to the standard A4 sized paper. The artifact of the ruler then serves as a way to save and carry forward a solution to the problem of measurement. In this way it also helps distribute the cognitive problem of measurement across representations, in some cases accruing social meaning. The solution may be more specific, for example, in the specialized artifacts for navigation, such as the Mercator projection [19], used to solve navigation problems.

Both the ruler and the map show one way that artifacts, which save and carry forward a partial solution, through a particular representation, can affect the process of an event far removed in time. To be successful however, such artifacts must be knit together with the mechanisms of current social practice. In the case of a well-defined problem like navigation, where the parameters that are important have been well understood for over a century, the social practice has been routinized and carried forward with the artifact. We believe organizational memory can be usefully viewed this way.

3.2. Analytical constructs used previously

Work by Star [35,36] notes that information can allow common activity between organizational groups when the information, represented in what she called a *boundary object,* provides enough shared understanding so that both groups have what they need. This can be extended to boundary objects across organizations, as in King and Star [22]. In both cases, only one of the groups has the full context for the information. For example, they may know the messiness of an employee's employment history and his associated records. The other groups using the record know merely that a person is employed at the company. The other parties do not have access to the full context; they are not overwhelmed by extraneous details (that is, extraneous to their needs). [5,6] detailed examples where simple statements (such as an employee's job title and dates of employment) can serve as validating employment to members of other organizations in a manner sufficient for their needs.

Mixed provenance [5-7] refers to the multiple origins and the associated validity and interpretation of organizational information. For example, a workflow system might have uses at the individual, group, and organizational level simultaneously. Moreover, provenance is one of several factors that help contextualize and interpret data. In our previous work we noted how hotline personnel choose different database sources for the same information based on their understanding of the validity and organizational issues surrounding the creation and maintenance of those databases. That is, users' perception of the provenance of each database affected their choice in any given search.

Finally *trajectories of use* [19,38] refers to the understanding that must happen for the information to be stored and appropriately reused. One must understand how the information will be used in some future in order to store useful information (or even to store any information). This includes the cycle of de-contextualization to remove extraneous details and the recontextualization used to make the information appropriate for the current situation. In one of the cases examined in [6], important information was "lost" because it was deemed sufficiently unimportant to be recorded. When a hotline representative tried to reuse the information in a call record, she was unable to properly decontextualize the stored record without appealing to the individual who handled the original call. In this sense, it can often be the traces of who was there that is important, not just the content [24,25].

While this work constructed the beginnings of analytical language for describing organizational memory it left many questions unanswered—particularly about the origins of memory stores and artifacts.

3.3. Artifact embodied expertise and reuse

These questions led us to consider another issue as well, how expertise and information reuse are related. Information can be sought many places, including information systems or people. One may think of expertise location in terms of knowing what we need to find out and looking for a person who can satisfy that need. But finding expertise takes many forms, and has many temporal extents. Seeking expertise and reusing information are both ways of information seeking.

Indeed, expertise can even be embedded in an artifact and an associated practice. For example, Hutchins [19] details navigation practice in a way that exposes the historical trajectory of the artifacts, as well as the training necessary to use those artifacts in various navigational processes. In this sense the artifact becomes the partially saved solution of a problem that has been solved many times before. Learning navigation becomes learning how to use the appropriate tools integrated into the process. However, it does not necessarily require understanding how to solve the problem from scratch. The tools themselves are the residua of the cognitive processes of an expert. Using the tools with the associated process has compartmentalized the problem in a way that no longer requires an expertise search for the Chief Navigator except in extraordinary conditions. One thing that makes this possible is that how to solve the problem is known and that formalization of the problem solution is acceptable both in navigation and in the military (where Hutchins' observations occurred). Thus, contextualizing the artifact (which is necessary for its use) has been carried forward with it in time. Its use is instantiated by distributing the problem solution between pre-saved solutions in the representation, the surrounding work process, and a variety of organizational and social roles.

In this paper we examine the life cycle of an artifact that was the outgrowth of a search for expert information, and which in turn became a resource in the environment. The genesis of the artifact is close enough in local memory that details of the original search that created it are still available. The artifact trace begins with a single sheet of paper, officially titled the "Summary of Departure Push and Arrival Rushes" and commonly known as the rush cheat sheet (RCS) at the Dallas TRACON (Tower Radar Approach Control).

Historically the RCS grew out of the process of analyzing traffic flow into and out of the Dallas-Ft. Worth (DFW) airport when they first decided to manage their traffic. Over time its form, status, and use have developed in both formal and informal ways. It can, and is, accessed and used as a resource informally, or as a prop for other information requests. It is also embedded in a variety of processes. The excerpts of current use will examine these changes in detail.

4. Study Site and Data

This study took place in two air traffic facilities. In the US, the Federal Aviation Administration (FAA) controls air traffic across the country with a tiered system of facilities that have differing geographical and conceptual responsibilities. *Towers* handle the area immediately around the airport and up to about 8000 feet, and rely largely on visual contact with the aircraft. *TRACONs* (Terminal Radar Approach CONtrol) are radar-based facilities that handle traffic in the terminal or airport area at approximately a 40-mile diameter and up to 20,000 feet. Radar based *Enroute Centers* handle the remainder, keeping traffic moving across the country.

The first author spent one year in the field observing in both the Dallas Ft. Worth TRACON (DFW) as well as in the Dallas Ft. Worth Center (ZFW). Data collected included detailed observations spread over 7 days per week and 24 hours per day at both facilities in order to catch seasonal, daily, and hourly variations. Data also included formal interviews of all four of the traffic management coordinators, their supervisor (and founding member), as well as informal interviews of many other personnel, including current controllers, supervisors, administrators, and retired controllers. Field notes were all transcribed and verified with the field site, as were formal interviews. Additionally the first author attended classes in being a traffic management coordinator at ZFW, and she was given access to internal documentation, software, and communications between the facilities for the process of traffic management. This data was supplemented by field notes from riding in the cockpit of numerous aircraft flying in and out of DFW in order to get the pilots' perspective. Data analysis was guided by distributed cognition theory. It relied on the numerous resource materials as well as ethnographic field notes and interviews.

5. What rush? The problem of managing traffic

Before describing the artifact and its history, it is important to provide an overview of the domain. That means understanding what a *rush* is in air traffic control (ATC) and why ATC is concerned with the problem of traffic management. The common sense notion that leaps to mind—rush hour traffic in any major metropolitan area—is mostly correct. In some cases, the traffic is direction specific. Everyone is going to the city in the morning or away from the city at night. Sometimes there are local conditions, an accident or construction, which make traffic difficult in both directions.

Just like car travel, air traffic has periods of concentrated traffic that happen throughout the day. These can be for arrivals, departures, or a combination of both. Airline schedules tend to cluster at certain times due to travel demand and competition, which results in peak times for arrivals and departures, or rushes[1]. However, there are differences between ATC and driving. First, one of the jobs of ATC is to try and keep traffic moving as smoothly as possible across the country. Second, they can direct or control traffic in order to ensure those smooth operations.

Prior to the founding of an integrated Traffic Management System (TMS) each facility operated individually taking traffic on a first come-first served basis. Using radar, each facility was only able to look a half an hour or so outside their boundary. This meant that problems elsewhere could impact another facility across the country with little or no warning. When traffic became too congested, aircraft were put in holding patterns, sometimes for hours at a time. The oil crisis in the 1970's made this solution unfeasible and sparked the founding of a nationally integrated system.

The Central Flow Control Facility (CF^2) collects information from local Traffic Management Units (TMUs) primarily in Centers across the country. The CF^2's goal is to determine areas of traffic saturation before they occur and then work with the local and other facilities to mitigate the effect elsewhere. Anyone in the US who has been told that "air traffic control has delayed our departure because of fog in San Francisco" is experiencing the result of a traffic management decision that has propagated nationally. Because fog at San Francisco forces the use of only one runway, the number of aircraft that can land per hour is almost cut in half. This means that the rushes that congest the airport must be evened out so that an aircraft from New York does not circle over San Francisco for an hour waiting to land, or worse, run out of fuel and have to be diverted.

Delaying aircraft on the ground at departure may keep them from circling at their destination hours from now and thus aggravating a difficult situation. However, it can create problems at the airport where the departure currently is, and it does not help solve the original local problem. As traffic congestion at individual airports has gotten more intense, the development of traffic management on a more local scale has occurred. This led to facility-based traffic management coordinators (TMCs) who have two jobs. One is to expedite the flow of local traffic. Locally, in order to decide how to handle a particular rush, the TMCs need information about rush composition (such as number of aircraft in a specific period, their type, as well as the direction of departure or arrival). The regularity and repetition in airline schedules means that TMCs have a good idea of what should happen, and they can adjust based on consultation with CF^2 about how things are unfolding on a particular day. Fog in San Francisco, thunderstorms in Dallas, or snow in Chicago can (and often does) ripple through airline schedules across the entire country, and TMCs must adjust. Local TMCs are also responsible for passing on appropriate information to the national level to help mitigate the effect local problems have on the nation.

While traffic management is itself an interesting problem [14] our interest is in understanding how one facility—the Dallas-Ft. Worth Terminal Radar Approach Control (TRACON)—started managing their local traffic. We are interested because this story is a search for expert information and the subsequent codification into a document used for a variety of purposes.

5.1. Traffic Management Artifacts

In the late 1980s when Dallas Ft. Worth (DFW) TRACON decided to create a traffic management unit (TMU), they assigned a controller[2] with a decade of experience at DFW to examine the problem. Ed Rondell[3] started out spending a week trying to get an overview of the flow of traffic. (Because each controller works only on a subset of a facility's airspace, no one controller would have had all the information about the traffic composition necessary to get the facility-wide view of the traffic required.) That done, he called around to the other TRACONs that managed their local traffic to find out how they "did it". What he found was that each of the three facilities (St. Louis, Chicago, and New York) did it differently. Some of those differences were related to local traffic and weather conditions, while others represent differences in each facility's control philosophy. Building on basic information about DFW's traffic and how other TRACONs' managed their traffic, he wrote a document that defined the philosophy of managing traffic at Dallas, and through that, the job of a local TMC. Then he set out to understand DFW's traffic.

Ed spent another year analyzing the flows into and out of DFW before he felt ready to begin to manage them. He started with a summary of traffic counts. As a matter of internal record keeping, arrivals and departures are recorded and broken down by hour (see left, Table 1.) (This count is important because a facility's rating and associated pay scales are determined by these totals.)

Table 1 shows the relative proportions of departures and arrivals. Rank ordering it by busiest to least busy hour helps to see when arrivals and departures are heaviest (as in the right hand side of Table 1). However, the raw numbers did not provide much of the information necessary for understanding why a situation occurs or what to do about it. For that kind of analysis he needed more detail within smaller blocks of time. He needed finer timing detail, the kind of aircraft involved, and the

[1] Departure rushes are also called pushes, but we'll use the single term.

[2] Two were assigned, but one was reassigned within a week.

[3] Not his real name.

Table 1. Hourly counts for arrivals and departures at DFW TRACON as of March 1993. The left hand side is ordered by time while the right hand side is ordered by total traffic. The busiest two hours are highlighted in bold.

Hour (LCL)		Arrivals	Departures	Total
0500	0600	19	6	25
0600	0700	5	26	31
0700	0800	60	42	102
0800	0900	77	88	165
0900	1000	58	71	129
1000	1100	54	33	87
1100	1200	84	88	172
1200	1300	53	55	108
1300	1400	74	96	170
1400	1500	84	49	133
1500	1600	49	52	101
1600	1700	72	71	143
1700	**1800**	**85**	**105**	**190**
1800	1900	74	83	157
1900	2000	65	35	100
2000	**2100**	**82**	**112**	**194**
2100	2200	57	14	71
2200	2300	22	70	92
2300	0000	10	12	22

Rank	Hour(LCL)		Arrivals	Departures	Total
1	**2000**	**2100**	**82**	**112**	**194**
2	**1700**	**1800**	**85**	**105**	**190**
3	1100	1200	84	88	172
4	1300	1400	74	96	170
5	0800	0900	77	88	165
6	1800	1900	74	83	157
7	1600	1700	72	71	143
8	1400	1500	84	49	133
9	0900	1000	58	71	129
10	1200	1300	53	55	108
11	0700	0800	60	42	102
12	1500	1600	49	52	101
13	1900	2000	65	35	100
14	2200	2300	22	70	92
15	1000	1100	54	33	87
16	2100	2200	57	14	71
17	0600	0700	5	26	31
18	0500	0600	19	6	25
19	2300	0000	10	12	22

direction of arrivals and departures to determine the rushes' start and stop times, as well as their composition.

For example, on the right hand side of Table 1 the busiest hour and the second busiest hour do not appear very different. Both have around 190 aircraft in an hour. The hour from 2000 to 2100 is the busiest for departures, while the second busiest hour (1700-1800) is the busiest for arrivals (see the highlighted rows in Table 1). However, the numbers do not tell the reader that the arrivals in the first case come in fairly steadily in the last 40 minutes of the hour, while departures fluctuate with more departures in the first 20 minutes. The second hour's arrivals build slowly for the first 10 minutes, then remain steady for 50 minutes, while departures peak in the first 30 minutes and then fall off steeply. Nor does the overview found in Table 1 indicate the relative proportion of aircraft or the direction of flow.

All of these details are aspects that affect what can be done to manage the traffic. For example, because Dallas is in the middle of the country, most aircraft come from either the east or the west, with fewer coming from the south and north. (In contrast, Chicago gets almost no traffic from the north.) Aircraft type determines the spacing between aircraft on departure. A smaller turbo prop needs additional spacing behind a heavy jet such as a 747 because of the larger turbulence caused during takeoff. Similarly, it also affects what aircraft are capable of on approach and therefore how they can be spaced.

This level of detail was not something that was being compiled in any one place when Ed began. They had daily numbers of departures and arrivals. With some additional work they could be split into time frames (as in Table 1), and with more work aircraft categories (turbo or jet) could be compiled. However, it was difficult to understand from a table of numbers what the nuances were from hour to hour. Furthermore, it was impossible to see what was supposed to have happened and what intervened on any particular day.

Ed turned to the airlines to begin to get a clearer picture of the traffic composition. Airline schedules gave an idea of expected times of peak traffic, as well as information about scheduled aircraft type. Additional information about where the traffic was funneled into the TRACON's airspace came from both his local facility, DFW, and the adjacent center, ZFW. Combining information from these sources he developed a document that detailed the hourly composition of the traffic. This pamphlet became known as the DFW Game Plan.

The Game Plan (GP) is an official, as well as a working document. It begins with a memo from the head of the TRACON outlining its purpose and use.

> "The Game Plan is a dynamic document to be used as a "guide" to provide an organized and effective management of traffic at DFW. It is based upon optimum conditions; supervisors and traffic management coordinators are expected to formulate alternate procedures as conditions warrant." [1]

Like other forms of organizational memory, use of the GP requires recontextualizing it to the current non-optimum conditions. To this end the GP consists of four different representations. First is the table with hourly traffic counts (as in Table 1), rank ordered with the busiest hour at the top. The second section is the "Summary of DFW Dept Pushes & Arrival Rushes", also known as the Rush Cheat Sheet (RCS), which we will discuss below. The third section further details the what is synopsized in these two one-page artifacts (Figure 1).

Each page shows two graphs covering arrivals and departures over a two-hour period. (Arrivals are on top and departures at the bottom of each page.) Pages overlap by an hour so a reader can get a better feel for the traffic. Thus the graphs for 0400 to 0600 local are on the facing page to the graphs for 0500 to 0700 local.

Figure 1 shows one page from this section (0600-0800), including the first morning rush between 7 and 8 am local time. Each graph splits the time into 10 minute increments. The different fill patterns delineate jet from turbo-prop aircraft so a reader can see at a glance the relative proportions of the aircraft, and thus infer the relative impact of their capabilities.

For busy periods like this example there are fairly extensive notes that detail what generally happens (the Synopsis), the direction the aircraft either arrive from or depart to (based on the points through which the aircraft are funneled), and the best way to handle the rush under "optimum" circumstances (the Scenario).

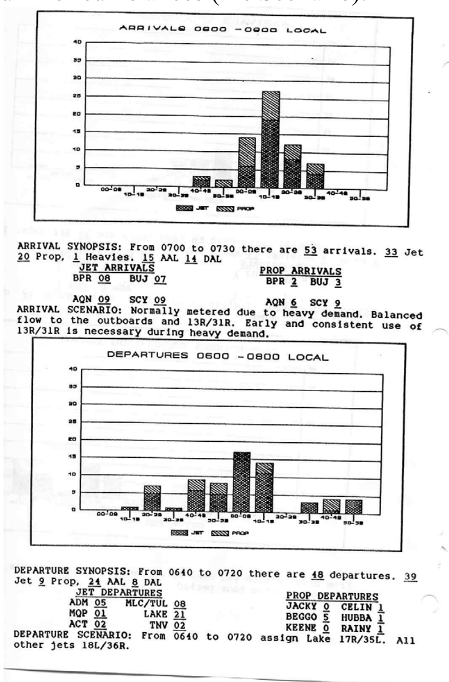

Figure 1. A page from the GP detailing 0600-0800 local time. Arrivals are at top, while departures are on the bottom graph.

Juxtaposing arrivals and departures in this way shows which sectors are going to get the traffic first, as well as how the airport itself will be impacted. For that reason this document is also used by the tower, and the fourth section is really for them: the Gridlock Prevention Plan. This section lays out, by direction of flow and predominance of departures or arrivals, which taxiways to use in order to avoid gridlock between arriving and departing aircraft. This is important in situations like that in Figure 1, when arrivals and departures are both strong.

Despite the fact that the GP is under 25 pages (usually printed double sided), it was too cumbersome to refer to on a daily basis. So, in addition, this document was reduced to a single schematic sheet, the RCS (see Figure 2). The RCS graphically depicts a synopsized version of the information in the graphs, with the lowest traffic times elided. (Note, while times are left off of Table 1 based on the FAA requirements for reporting, the times left off here are based on the TMCs preferences. For example, the first rush of the day, very early in the morning, is a cargo rush, and is not shown.)

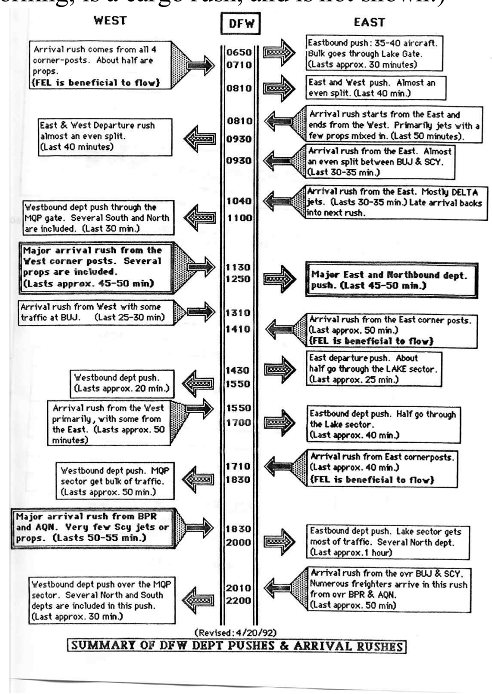

Figure 2. Rush Cheat Sheet. In the center is a time line that orients rushes to east and west.

The RCS schematizes the detail in the GP graphs, decontextualizing it somewhat to fit the restricted space but adding context in the representation. For example, the top three entries on the RCS cover the same period as the graphs in Figure 1. It shows the arrival rush from the west, bracketed by departures predominantly going east. However, much of the detail has been left out. What remains is the kind of information necessary for a quick reminder. Direction of flow is generalized to east and west around the timeline and comments such as "Arrival rush comes from all 4 corner-posts[4]." In contrast, Figure 1 shows exactly how many aircraft come over each corner-post. The GP also includes details about which airlines compose the rush, as well as the number of each

[4] Corner posts are navigational fixes on the edge of the TRACON's airspace that are used to sequence arrivals. Because of the orientation of DFW and its airspace these fixes form a box around DFW. Hence the name corner-post.

aircraft type. This last detail is abbreviated on the RCS to statements like “About half props.”

The letter size of the RCS meant it could be placed under the glass counter top of the TMCs’ work surface so that it was easily accessible—not only to them but to others. As the TRACON traffic management coordinators used this document, and displayed it in their work area, other members of the community began to refer to it as well.

6. Three uses internally, and externally?

The RCS provides a visual reminder of the traffic throughout the day. Its organization takes advantage of the way in which the DFW airport has runways mostly oriented north-south and has traffic predominantly arriving from the east and west. Because TMCs handle the traffic every day, they can generally rely on their own memory and expertise for the details instead of directly depending on the RCS. They often reason about, and in reference to, the RCS without looking at it. They use the RCS largely after an absence, whether a short break or a long vacation. In these cases the brevity of the rush characterizations is sufficient. With a controller’s training, a TMC can interpret and distinguish between a characterization that reads “Arrival rush comes from all 4 corner-posts,” and “Major arrival rush from the West corner posts”, and be able to answer queries from other facilities. (While they do not refer to it constantly, it does present the information in a way that is salient. Many times during observations, particular situations were explained to the first author with reference to the RCS.)

The RCS also began to be used by other AT personnel in unexpected ways. Controllers and supervisors refer to it as their own reminder, using it or references to it in order to negotiate breaks. A controller on break might pop back into the control room to see whether he has enough time for another cigarette. In asking his supervisor if it is ok they may both refer to the RCS either directly, or in talking to a TMC, checking against current radar data to confirm whether the schedule is on time or delayed, hence our title.

Facility administrators also use it to decide when might be a good time for a tour. (This is perhaps less surprising when you know that administrators started out as controllers.) Again, the RCS becomes an artifact around which people discuss and reason. Several times during observations administrators came into the control room to check out good times for a tour. This was defined by when there would be enough traffic to be interesting, but not so much that it would either be incomprehensible, or the tour would be intrusive on operations. In this they rely on the schematic of the RCS combined with their own knowledge of the facility with additional guidance from the TMC. Most times TMCs responded to such requests with a minimal response, only discussing to point out that an inference about the character of a particular rush was incorrect.

The GP is used by others as well. As the GP was being formalized and reduced to the RCS, the airlines realized that it could help them reason about their business. While the airlines are the source of much of the GP’s information, they do not have the complete picture.

The synthesis across different companies, incorporating controller experience, means that the GP gives the airlines a more accurate picture of the variation in airport load over the day than they could have otherwise. It includes detail about the number of aircraft from the two busiest airlines at DFW, as well as the more variable load from cargo and private planes. Further, the synthesis with local ATC knowledge means they can understand why certain aircraft always run behind schedule. For example, an aircraft’s placement late in an arrival rush from the west may mean it is almost always diverted over the airport to come in from the east, resulting in both longer flight times and taxi times. Moving the schedule up or back as little as 15 minutes could make the difference in being on time—not only here but as that aircraft’s individual schedule ripples across the county. At least one airline has used such information in decisions about schedule changes.

7. Discussion

In the RCS and the GP we have a set of data embedded in artifacts that work in a number of ways related to memory–both individual and organizational. Tracing their development exposes a number of issues that clarify our earlier analytical constructs.

7.1. Formalizing the representation of solutions

This case shows the role formalization can have in allowing an organization to adopt (and adapt) an artifact as a memory resource. To a large extent, this is similar to Bannon and Bodker’s [9] punctuation in informational artifacts, moments when informational artifacts cease to be dynamic and changing, and instead crystallize. Here, however, we note that the RCS and GP continue to be dynamic in content, but the form must crystallize. We believe this to be true for most boundary objects.

The RCS synthesizes data across a number of dimensions. References, like **{FEL is beneficial to flow}** (where FEL is the name of an alternative to a corner post for handling heavy traffic) or **Use Lake for departures** (see Figure 2), codify past experience in a way that is visible and comprehensible at a glance for the TMCs, though obscure to outsiders. These excerpted statements represent an accepted solution to repeating traffic problems, which have been thus saved and formalized in the RCS. As a formalization, it can be adopted by the organization as a memory resource and then used.

For example, in the early morning rush the notation **{FEL is beneficial to flow}** refers to using the additional fix FEL to help feed the traffic. A TMC, noticing that the Center is not feeding traffic through the additional fix FEL, will call them to find out why. In the course of normal work, however, this solution to spreading out the morning rush will be so routinized at both facilities that nothing will need to be said. In fact, the pilot of an aircraft regularly placed in the rush so that it is necessary to use FEL will expect it to be assigned. She will likely be prepared with the appropriate arrival charts opened, based on her previous personal experience.

This formalization may occur because air traffic is very repetitive and therefore its control can be routinized. Routinization is fostered because safety is paramount. ATC regularly formalizes individual solutions as they are found to work repeatedly. Once such a solution is found and validated (a nationally and locally defined process), formalization ensures accuracy and therefore higher safety of operations. However, organizations are full of similar examples of formalization from repetitive action.

The RCS itself codifies information related to the nature of the rush (departure or arrival) and its composition (aircraft type—jet or turbo—and number) to the basic orientation of the airport and the time of day. This representation, as we saw earlier in Hutchins' navigation example, is a saved partial solution. TMCs know that direction is important to understanding the nature of the traffic flow and what can be done to handle it. It may be serendipitous that the runway configuration is so easily transformed to the presumed "north is up" on a page; however, the result is powerful. Similarly, the representation of the timeline overlaid on top of the directions can be used quickly to orient use throughout the day. Interestingly, while most of ATC is done based on GMT, these documents are all in local time. These representational choices make it easy for a newcomer to quickly orient to and use the document.

Over time DFW TRACON has settled on three representations that work for different purposes, through a process of evolutionary adoption and adaptation. The basic count and rank ordering is forwarded nationally to document their traffic level. The GP and its integration of cross-organizational information are primarily useful to the airlines. Finally, the RCS is maintained at the TMC position and is used in a variety of ways. Note that although only the RCS is useful internally, it evolved from the basic count and rank ordering and the GP as the problem became better understood. In a sense, its form could only be discovered as other representations were used to solve the problem. The basic count and rank ordering and GP have continued to exist as boundary objects, indeed almost entirely for the use of others.

7.2. Boundary object abstraction

Boundary objects, in general, remove detail from the purview of external users to avoid drowning them in details. In this case, however, the abstraction is reversed. The reduction is done for internal users; the RCS has noticeably less detail than the GP. On the other hand, the GP must have additional detail in order to be useful to the airlines as a boundary object. It will be necessary for further studies to understand more generally when this reversal is required in boundary objects.

7.3. Mixed provenance and information validity

Earlier work noted that boundary objects require authoritativeness and veracity in use. This case confirms that finding, but adds a twist. Not only are the RCS and the GP the outcome of a synthesis of data across a number of dimensions, the synthesis also occurs from several organizations. Data from multiple airlines and cargo companies is further integrated with local facility experience drawing on weather patterns, local conditions, and control philosophies of adjacent facilities. While the data has mixed provenance (coming from different sources), its validity is largely unquestioned.

The accuracy and validity of the information is assumed, perhaps because the FAA acts as the point of collection and synthesis. Safety is paramount in the air transport system [14,24], and if the data are released, members of the air transportation industry will assume its validity. This is in interesting contrast to hotline work [5-7], where telephone agents were concerned with the degree of degradation in data. Here, the boundary object is either valid and released, or non-existent.

Unlike our hotline study, the source of the information used, whether the RCS or GP, is not based on variations in presumed validity or context. Both have the same authority, so use is more determined by a question of scale. While the GP is an analytic precursor to both the creation and ongoing updates of the RCS, it is no longer referred to by TMCs to support the details of their work. The RCS reduces the detail necessary in the GP for understanding the traffic to a representation that is more suited for managing the traffic.

The data's validity, along with the shared knowledge and language of the domain, mean that while the data underlying the RCS and the GP can act as a resource for both the airlines and ATC. However, the preferred representation varies. The airlines want the detail of the GP while TMCs and other AT personnel prefer the RCS.

The individual trajectory of use here is much less important. In some sense the document, and the information collected for it, is so specialized to the task that it is a perfect match to the expected use. Thus the risk of missing information is reduced. The stability of the form means that the data is mostly in the appropriate context, so recontextualizing the data is similarly less problematic. While the airline may not use the data in the same way a controller does, the data is presented in sufficiently comprehensible terms that it can be used in reasoning without needing some detailed translation.

Critical information is rarely missing, thus, finding the individual who provided a piece of data is unnecessary. Such searches only tend to occur when the RCS and GP are updated. Clarification may occur between the organizations, but they tend to happen after the moment.

7.4. Data Maintenance

At this point, from the perspective of the TMCs, they could skip the process of creating a GP and make any needed amendments directly to the RCS. Amendments may be driven by changes to the airport or configuration of the airspace. More usually, though, the changes are in response to changing airline schedules. Compared to the job of originally compiling the GP and working out traffic management solutions for the facility, such updates are relatively minor. Most changes are the result of airline and cargo schedule changes—such as moving a flight back an hour or adding a new route[5]. These changes do not materially change the existence of a rush, although they might alter the duration. Thus, most changes would require minor alteration of the RCS.

The routinized nature of the overall role of traffic management at DFW, coupled with the repetition of the daily traffic pattern, means that the RCS as a memory artifact is now mostly referred to by people who do not handle the details of traffic management as part of their daily work. Those who do—the TMCs—maintain the knowledge internally. However, when they do need it the stability of the form helps them continue to use the document. (Minor changes to the RCS are still released, drawing attention to the dynamics of the data).

In contrast, the level of detail in the GP requires more detailed changes to keep it up-to-date. That level of detail is no longer necessary for the day-to-day work of managing the traffic at DFW, and so TMCs would like to stop maintaining the GP[6]. However, there may be a secondary benefit from its maintenance. The GP has become a point of communication between the airlines and ATC. It gives everyone a common representation around which they can negotiate the differing goals for their organizations. In this sense it may help them to create a common information space [9]. Understanding this fully will require further observations on the side of the airlines and their interactions with ATC.

It is important to note that the term “boundary object” implies a static, unmoving thing. Instead, we observed how critical evolutionary adaptation and adoption are at the macro-level and how critical regular updating of the data is at the micro-level. Indeed, it may be more appropriate to speak of *boundary streams* than objects.

8. Conclusion

The Rush Cheat Sheet is emblematic of many organizational memory artifacts. Its origin in the Game Plan means that it is driven by two organizations: the FAA and the airlines. While the details are dependent on the specific domain we have identified four general aspects that come together to make these artifacts work within and across organizations:

First, a solution must become crystallized in order to become a memory resource in an organization. Repetition of a problem drives the subsequent representation of the solution in a formalized manner. To be useful (and perhaps to be adopted as a resource), this solution must be maintained over time. Once determined useful the form remains relatively stable for organizational processes to work. However, the data must be dynamic especially to act across boundaries. We suggested, accordingly, that *boundary streams* might be more appropriate as a term.

Second, the underlying data changes so the forms must be maintained so the data can be adequately used. In a restricted case like ATC the problem itself is stable enough that discovery of the appropriate form, such as the RCS, will likely remain stable for quite some time. An underlying process is maintained to formalize data upkeep. In less regularized settings, such as offices, re-contextualization is necessary (e.g. a hotline record) because the problem and its solution changes, even when it can be described with similar words.

Third, while most boundary objects work by eliding unnecessary information, here the essential composition of the information relies on cross-organizational sources. While details are still elided, however these artifacts are only useful because of inter-organizational information.

Finally, sharing information is only useful if the information is perceived to be true. Unlike most office settings, air traffic control and air transport in general, is a domain where safety is paramount. In contrast to the political and technical reasons that make for competing and questionable data in the workplace, in ATC data simply is not available unless it has been validated. This argues that choices of artifact are not based only on validity as we have seen before, but on the utility of the representation for the task at hand.

Using distributed cognition theory as lenses focusing our analysis towards representations and processes helped expose their interplay in the development of an artifact for information reuse. Continuing explorations of such naturally occurring examples will lead us to a better understanding of which concepts are generalizable and which are limited by the domain.

[5] The first author has a copy of the 1993 edition , in use at the time of the observations. At that time the RCS had been amended 39 times since its origin (late 1980s). (There was disagreement among TMCs whether the RCS originated in 1986 or 1988.)

[6] Note however, that adding a new runway, (as was done in Dallas in the late 1980s) has a much more extensive impact, causing alterations in airspace, departure and arrival routing and taxiway routing which in turn would likely require changes at the level of detail in the GP.

Acknowledgements

This project has been funded, in part, by NSF grant IRI-0124878. We thank Ed Hutchins, Wayne Lutters, and Dave McDonald for useful discussion. We also thank the air traffic personnel at DFW and ZFW for their enthusiasm, patience and insights. The data was collected under NASA grant NAG2-873 and benefited from the support of Tom Davis and Len Tobias.

9. References

1. , Dallas-Fort Worth Game Plan. Dept. of Transportation, FAA, Spring 1993 edition.
2. Ackerman, M. S. Answer Garden: A Tool for Growing Organizational Memory. MIT, Ph.D. Thesis, 1993.
3 Ackerman, M. S. Augmenting the Organizational Memory: A Field Study of Answer Garden. Proceedings of Conference on Computer-Supported Cooperative Work , 1994: 243-252.
4. Ackerman, M. S. Definitional and Contextual Issues in Organizational and Group Memories. Information Technology and People, 1996, 9(1): 10-24.
5. Ackerman, M. S. and C. Halverson. Considering an Organization's Memory. Proceedings of CSCW , 1998: 39-48.
6. Ackerman, M. S. and C. Halverson. Organizational Memory: Processes, Boundary Objects, and Trajectories. Proceedings of the IEEE Hawaii International Conference of System Sciences (HICSS'99), 1999:
7. Ackerman, M. S. and C. A. Halverson. Re-Examining Organizational Memory. C ACM, 2000, 43(1): 58-63.
8. Ackerman, M. S. and T. W. Malone. Answer Garden: A Tool for Growing Organizational Memory. Proceedings of ACM Conference on Office Information Systems, 1990: 31-39.
9. Bannon, L. and S. Bødker. Constructing Common Information Spaces. Proceedings of European Computer Supported Cooperative Work conference , 1997: 81-96.
10.Bannon, L. J. and K. Kuutti. Shifting Perspectives on Organizational Memory: From Storage to Active Remembering. Proceedings of Hawaii International Conference on System Sciences (HICSS-29), 1996: 156-167.
11.Chaplin, D. Community Memory. Department of Computer Science, Lancaster University, manuscript, 1994.
12.Clark, A. Being there: putting brain, body, and world together again. MIT Press, Cambridge, Mass, 1997.
13.Conklin, J. Corporate Memory. Proceedings of Groupware '92, 1992: 131-137.
14.Halverson, C. A. Inside the Cognitive Workplace: New Technology and Air Traffic Control.UCSD,Ph.D. Thesis, 1995.
15.Hollan, J. D., E. L. Hutchins and D. Kirsh. Distributed cognition: toward a new theoretical foundation for human-computer interaction research. JHCI 2000, 7(2): 174-196.
16.Huber, G. P. A Theory of the Effects of Advanced Information Technologies on Organizational Design, Intelligence, and Decision Making. Academy of Management Review, 1990, 15(1): 47-71.
17.Hughes, J. and V. King. Paperwork. University of Lancaster, Esprit COMIC Project, COMIC-LANCS-4-1, 1992.
18.Hughes, J., V. King, T. Rodden and H. Andersen. Moving Out from the Control Room: Ethnography in System Design. Proceedings of ACM Conference on Computer-Supported Cooperative Work, 1994: 429-439.
19.Hutchins, E. Cognition in the Wild. MIT, Cambridge, 1995.
20.Hutchins, E. and T. Klausen. Distributed cognition in an airline cockpit. In Middleton, D. and Y. Engestrom (ed). Communication and Cognition at Work. Sage Books CA.,1992.
21.Kamiya, K., M. Roscheisen and T. Winograd. Grassroots: A System Providing a Uniform Framework for Communicating, Structuring, Sharing Information, and Organizing People. Comp. Networks and ISDN Systems, 1996, 28(7-11): 1157-74.
22.King, J. L. and S. L. Star. Conceptual Foundations for the Development of Organizational Decision Support Systems. Proceedings of Hawaii International Conference on System Science, 1990: 143-151.
23.Lindstaedt, S. N. Group Memories: A Knowledge Medium for Communities of Practice. Department of Computer Science, University of Colorado, Ph.D. proposal, 1996.
24.Lutters, W. G. and M. S. Ackerman. Achieving Safety: A Field Study of Boundary Objects in Aircraft Technical Support. Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW'02), 2002: accepted for publication.
25.McDonald, D. W. and M. S. Ackerman. Just Talk to Me: A Field Study of Expertise Location. Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '98), 1998: 315-324.
26.Morrison, J. and M. Weiser. A Research Framework for Empirical Studies in Organizational Memory. Proceedings of 29th Annual Hawaii International Conference on System Sciences, 1996: 178-187.
27.Norman, D. A. Cognitive Artifacts. In Carroll, J. M. (ed). Designing interaction. Psychology at the human-computer interface. Cambridge University Press, 1991. 17-38.
28.Orlikowski, W. J. Learning from Notes: Organizational Issues in Groupware Implementation. Proceedings of Computer Supported Cooperative Work (CSCW'92), 1992: 362-369.
29.Roberts, J. The self-management of cultures. In Goodenough, W. (ed). Explorations in Cultural Antrhopology: Essays in Honor of George Murdock. McGraw-Hill, 1964.
30.Robinson, M., M. Kovalainen and E. Auramaki. Diary as dialogue in papermill process control. 2000, 43(1): 65.
31.Sacks, M. On-the-Job Learning in the Software Industry: Corporate Culture and the Acquisition of Knowledge. Quorum Books, Westport, CT, 1994.
32.Salomon, G. (ed). Distributed cognitions. Cambridge University Press, Cambridge, 1993.
33.Sandoe, K. and L. Olfman. Anticipating the Mnemonic Shift: Organizational Remembering and Forgetting in 2001. Proceedings of International Conference on Information Systems (ICIS), 1992: 127-137.
34.Smith, J. B. Collective Intelligence in Computer-based Collaboration. Lawrence Erlbaum , Hillsdale, NJ, 1994.
35.Star, S. L. The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving. In Gasser, L. and M. Huhns (ed). Distributed Artificial Intelligence. Morgan Kaufmann, San Mateo, 1989. 37-54.
36.Star, S. L. and J. R. Griesemer. Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. 1989, 19: 387-420.
37.Stein, E. W. and V. Zwass. Actualizing Organizational Memory with Information Systems. Information Systems Research, 1995, 6(2): 85-117.
38.Strauss, A. Creating Sociological Awareness: Collective Images and Symbolic Representations. Transaction, New Brunswick, 1991.
39.Weiser, M. and J. Morrison. Capturing, Linking and Retrieving Team Project Information. Working paper, Dept of Management Sciences, University of Iowa, 1994.