An ecosystems approach to human service database design

Derek Coursen

Keywords

human services, systems analysis and design, data models, client-tracking systems, case management systems, ecosystems perspective, ecomaps, knowledge management, performance measurement

Abstract

Early client-tracking databases were strongly influenced by the structure of previous paper-based systems. More recently, there has been evolution toward databases that attempt to represent the interconnectedness of people in the human service environment. No consensus on best practices, however, has yet emerged. This paper presents a systems analysis technique and a data model based on one of the theoretical foundations of current social work practice: the ecosystems perspective. This approach facilitates a range of knowledge management and performance measurement capabilities that have so far been uncommon in client-tracking systems.

Author/Acknowledgements

Derek Coursen is Director of Information Management at the Vera Institute of Justice. He has MS degrees in information and library science (Pratt Institute), information systems (Pace University) and management (New York University – Wagner School of Public Service). Address correspondence to: Derek Coursen, Vera Institute of Justice, 233 Broadway, 12th Floor, New York, NY 10279 (E-mail: dcoursen@vera.org).

I thank Jean Callahan, Evan Elkin, Krista Larson and David Ries for pilot testing this approach in Adolescent Portable Therapy; Quinn Cushing, Connie Knapp, Jim McMillan and Judith Ullman for reviewing the data model; Inna Itskova for turning design into software; Chris Stone for encouraging exploration of these ideas; and Megan Golden for providing the leadership and support without which this approach could not have been developed.

Introduction

Human service programs of all kinds keep records on their clients and the services provided. Indeed, case recording practices are recognized as one of the basic architectural elements of human service programs (Lohmann & Lohmann, 2002). At one time, case records were universally kept in paper files. Recent decades, however, have seen the rise of computerized client-tracking systems. The replacement of paper by databases has led application developers to reexamine the structure of human service case information.

There is considerable uniformity across the sector regarding the structure of paper-based case records. Programmatic differences notwithstanding, the common approach revolves around the notion of a *client chart*, the traditional multi-leaved manila folder format. The client chart includes a fact sheet with the client's name, address and phone numbers, family information, medical issues, and other basics; a set of data on the client's history and needs, which is collected during the intake interview; and a log of case notes in which staff document their service activities related to the client (Rosenthal, 2003). These and other elements—such as the results of clinical tests and information on program discharge—are kept within a chart which is identified as pertaining to a single index client.

The client chart is a useful format because it arranges a considerable body of information in a way convenient to human service workers. Program staff know where in a chart to find or amend certain relatively stable data (e.g. date of birth, current address) and where to add or find new information about the client's progress. Because the chart is so useful, it has informed the development of client-tracking software, where it is common to find such elements as a fact sheet screen, a case notes screen, and a client chart report. In short, the conventions of the paper file have been reproduced in software interfaces.

The influence of the client chart extends, however, beyond interfaces. Rather than merely guiding presentation, the client chart has often molded the underlying structure of the information. This occurred because *semantic conventions*—commonly accepted ways of modeling particular kinds of situations—are the foundation of data models (Hay, 1996). The client chart provided simple and seemingly intuitive semantic conventions which shaped the data models of the first generation of these systems.

As human service software has matured, the limitations of data models based on the client chart have become apparent. One fundamental problem is how to represent interconnectedness: the fact that the human service environment is a complex web of connections and interactions in which the same people may appear in different

roles in different cases over time. Software developers have begun to address this issue but a consensus on best practices has yet to emerge.

In this paper, I will present an approach to the problem based on the ecosystems perspective. The approach begins with a technique for analyzing a human service environment as a "programmatic ecosystem." The structure of that analysis directly corresponds to a set of semantic conventions leading to a data model. Analysis of any given program's ecosystem thus allows the data model to be applied in that situation.

A client-tracking system based on this model can offer a range of capabilities which are not yet widespread in human service information systems. These include sharing information on the interconnections of people and organizations across different cases, and delivering performance indicators on the involvement of non-client constituents.

Data models in client-tracking systems

There has been little formal literature on human service data models. Though theorists have created detailed handbooks of the data models patterns found in commerce, accounting, insurance, manufacturing, telecommunications, and other industries (Hay, 1996; Silverston, 2001a, 2001b), these resources do not specifically cover the human services. Searches in Social Service Abstracts and Science Direct reveal sparse work on the subject. Practitioners who have designed client-tracking databases or evaluated commercial ones will, nonetheless, recognize patterns that commonly occur.

Typically, in early versions of client-tracking systems, the core of the data model is an independent entity representing the client. There is a one-to-many relationship between the CLIENT entity and a dependent entity representing the service activities related to the client. This SERVICE EVENT entity contains attributes representing the date, type of service, case note, and the staff member who provided the service. The data model also includes components that represent non-client constituents such as the client's family members and significant service providers (e.g. physician, attorney, parole officer). In very immature models these may be implemented as attributes in the CLIENT entity, but it is more common to find a one-to-many relationship between the CLIENT entity and one or more dependent entity representing these people. Figure 1 shows an example of a data model that has these patterns.

Such a data model is based on semantic conventions that can be summed up as follows: "The program serves clients; the program may provide multiple service events to each client; each client may have family members and external service providers." While these statements are correct as far as they go, they are clearly derived from the structure of a paper-based system, where there is one chart per client case, the chart has a log of service activities, and there are places in the chart to note other important people in the client's life.

This model has severe limitations, the most evident of which is in the treatment of people. Modeling non-client constituents as dependent upon the client fails to account for the fact that people may show up in different roles at different times. This makes it impossible for the program to maintain a unified store of data on a particular individual.

FIGURE 1. Client Chart Data Model

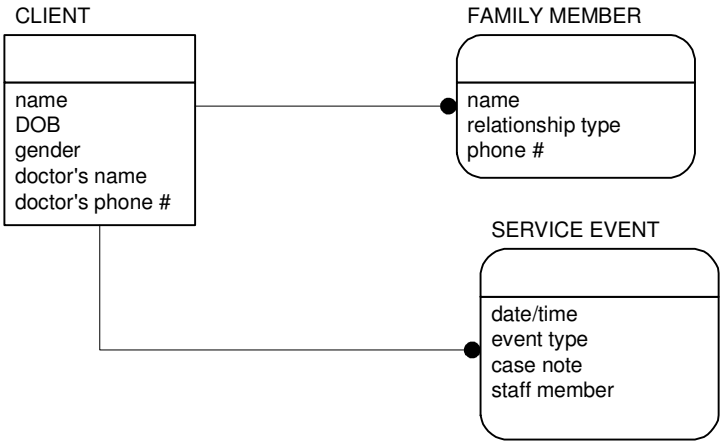

A glance at recent software packages developed in various environments—in-house, commercial, and open-source—will show that developers are increasingly aware of this issue. Efforts to address it often include recognition of the fact that there are various roles that people may play vis-à-vis the agency, and that people may be connected to each other in family or household arrangements.

A particularly intentional and well-documented example comes from the IT Resource Center (ITRC). That Chicago nonprofit organization builds software for other agencies and has publicly shared its own database design to encourage exploration of these issues. (For consistency in comparing different models, I will refer to some entities in the ITRC's logical data model by names slightly different from the published table names.)

Like the client chart data model presented above, the ITRC model (Mills-Groninger, 2000) has a CLIENT entity in a one-to-many relationship with a SERVICE EVENT entity. However, the CLIENT is in fact a subtype—though the only one—of the supertype INDIVIDUAL, which represents any person that the agency may need to track. The INDIVIDUAL has a one-to-many relationship with a PERSON TYPE entity that indicates how the individual is related to the agency (e.g. as a client, staff member, external professional, etc.). The INDIVIDUAL also has a many-to-many relationship with a HOUSEHOLD entity. And in some implementations, the data model has been extended to include a structure entity FAMILY RELATIONSHIP linking the INDIVIDUAL with itself (T. Mills-Groninger, personal communication, August 23, 2004).

The semantic conventions underlying this part of the ITRC data model, then, may be stated as follows: "The agency has clients, and may provide multiple service events to each client. However, the agency relates to many individuals in many ways; an individual being a client is just one of those ways, and the same individual may be related to the agency in more than one way at any given time or over time. There are households comprising multiple individuals, and an individual may belong to one or more household. Individuals may have family relationships with each other."

These statements are considerably more comprehensive and generic than those based on the client chart. Data models like the ITRC one are, therefore, better able to track the universe of people in a human service environment. Agencies using the ITRC's software have benefited from its ability to represent the fact that, say, a former client of the agency has become a supervisor capable of hiring current clients; the software thus helps users see some of the interconnectedness that exists in the real world (Mills-Groninger, 2003).

There are, however, even deeper levels of interconnectedness that can be tracked. And human service programs, by their nature, have an interest in interconnectedness. In fact, it is central to current social work thinking, where a theoretical framework based on it—the ecosystems perspective—has been widely accepted in recent decades. That perspective is relevant to database design because it offers an approach for advancing even further in understanding and representing interconnectedness.

The ecosystems perspective

The ecosystems perspective draws on insights from ecology and general systems theory. Much of the perspective was shaped by the work of the anthropologist Gregory Bateson, whose ideas entered the social work field via their influence on psychiatry and family therapy (Bilson & Ross, 1999).

According to Meyer and Mattaini's (1998) presentation, the ecosystems perspective is a way of approaching a case as a complex system of interconnected phenomena, and of considering the client's interactions with multiple factors and actors in his or her environment. By highlighting the interacting elements in a case, the approach is intended to provide the practitioner with multiple possibilities for intervention. In doing so, it reinforces the distinction between social work and those disciplines which focus on individual dysfunction. In this perspective, the individual is seen as both adapting to the environment and affecting it. The individual operates within a system that has boundaries; elements of the system interact with each other reciprocally; and the system as a whole tends to seek a steady state. The notion of equifinality holds that interventions at different points in the system may conduce to the same final effect. The converse, multifinality, notes that a single event has multiple effects in different parts of the system. In practical terms, the approach suggests that the social worker consider a range of interventive options and act on those points in the client's ecosystem where there is an opportunity for positive change.

Practitioners often use a visual formalism called an *ecomap* (Hartman, 1978) to conceptualize the ecosystem of an individual case. A simple ecomap will show the important people and institutions that interact with the client. More elaborate ones include activities, available and missing resources, and positive and aversive influences.

Ecomaps may also be used to represent the common environment of an entire client population. Mattaini (1993) suggests using *organizational* and *nomothetic* ecomaps. An organizational ecomap

shows how clients interact with the human service agency's subsystems and other external systems. A nomothetic ecomap shows the kinds of actors and factors that tend to interact with a group of cases in a community.

For the purpose of understanding interconnectedness in a human service program, a new variation on this idea will be useful. The ecosystem of a program can be analyzed in terms of the types of people and organizations in the environment, how those types relate to each other, and how the program attempts to intervene with them. The results of this analysis can be represented graphically as a "programmatic ecomap," the structure of which will directly correspond to a data model. This technique of linking a graphical analysis methodology to a data model convention has previously been found useful in another area of human service information management: modeling the complex ways that participants flow through programs (Coursen & Ferns, 2004).

Analysis of the programmatic ecosystem

A program's ecosystem can be analyzed in a sequence of three steps. Each step is an inventory of a different kind. Together, they provide an overview of the program's nature and the kinds of information that the program needs to capture.

The first step is an inventory of all of the types of people and organizations (or "parties" for short) that exist within the program's ecosystem, i.e. with which both the clients and the program interact. The types are defined based on typical roles that people or organizations may play vis-à-vis the program, such as client, family member, external service provider organization, etc. Definitions therefore reflect the program's philosophy and operating environment.

The second step is to determine, for this set of party types, what types of significant recordable relationships typically exist among them. In this context a "relationship" means a formally definable connection that exists for some period of time, such as a family, caseload or employment relationship. Most types of relationships will involve program participants, but there are usually important ones that do not.

The third step is to list the types of events that the program needs to record. Most types of events—e.g. counseling sessions—are within the program's work. They are the ways the program attempts to intervene in the client's ecosystem. Some types, however, may be events outside the program—such as educational milestones or criminal arrests—which affect the client's life. The nature of each type of event typically implies a set of business rules about the parties of different types that must or may be involved. For example, an external case conference might be defined as requiring the participation of both a program staff member and a staff member from an external service provider.

The three inventories can then be represented graphically together as the programmatic ecomap. Shaded and unshaded ovals stand for organization types and person types, respectively. Thick lines between party types stand for relationship types. Event types are boxes, and are linked to party types with thin lines.

The following example illustrates analysis of the programmatic ecosystem in a real human service program.

Case example: Adolescent Portable Therapy

Adolescent Portable Therapy (APT) is a demonstration project of the Vera Institute of Justice, a nonprofit organization that develops innovations in justice system practices. APT is an intensive, outpatient, family-based service for the most serious drug abusers within New York City's juvenile justice system. The core innovation of APT is to deliver continuous substance abuse treatment to these youths no matter where they are. The program's therapists treat participants whether they are in pre-trial detention, in a juvenile corrections facility, at home or elsewhere.

During 2001, the program's first year of operation, Vera developed generic client-tracking software and customized it for APT. In the course of that customization, software developers and APT staff analyzed the programmatic ecosystem as follows:

The focus of the program is, of course, the *participants*. APT also works closely with *collaterals*—the key people in the participants' nuclear, extended and informal families. Treatment is delivered by *program staff*. Participants often live at *placement sites*, court-ordered residential situations. Depending on the nature and affiliation of the placement site, participants there may be supervised by *juvenile justice facility staff*, *child welfare staff* or *private agency staff*. As participants wend their way through the juvenile justice system, the program interacts with *lawyers*, with *probation officers*, and with *aftercare staff* who check up on participants following release. After participants return to their *schools*, APT works to promote academic success by keeping in touch with

teachers, *guidance counselors*, and *school administrators*. And APT staff refer participants to needed services and constructive recreational activities offered by *resource organization staff* at *resource organizations* in the community.

Participants are on the *caseloads* of program staff; externally, they are also on the caseloads of child welfare staff, probation officers and aftercare staff. They have various *family relationships* with collaterals. They are *placed with* placement sites. They are *represented by* lawyers. They are *enrolled with* schools. And at the schools, they are *assigned to* particular guidance counselors and teachers. Finally, because APT interacts both with external institutions and with individuals who represent them, connections between these people and organizations are programmatically important. Thus external organization staff of various types are *employed by* the placement sites, schools and resource organizations.

Program staff provide both *individual therapy sessions with participants* and *individual therapy sessions with caregivers*. As the name indicates, an individual therapy session involves one therapist and one other person. There are also *family therapy sessions* involving more than one member of the participant's family, usually but not always including the participant. Program staff make *referrals* of participants to resource organizations. And they have *external case conferences* with actors from the juvenile justice system, placement sites, schools and resource organizations.

Figure 2 is a programmatic ecomap of Adolescent Portable Therapy based on the party, relationship and event types mentioned above. (Because of limited graphical space, relationship type names and the precise business rules for events are omitted from the ecomap. And for the sake of brevity, not all details of APT's ecosystem have been described here. For example, there are many other event types representing other interventive techniques and formats, the administration of diagnostic instruments, etc.)

FIGURE 2. Programmatic Ecomap of APT

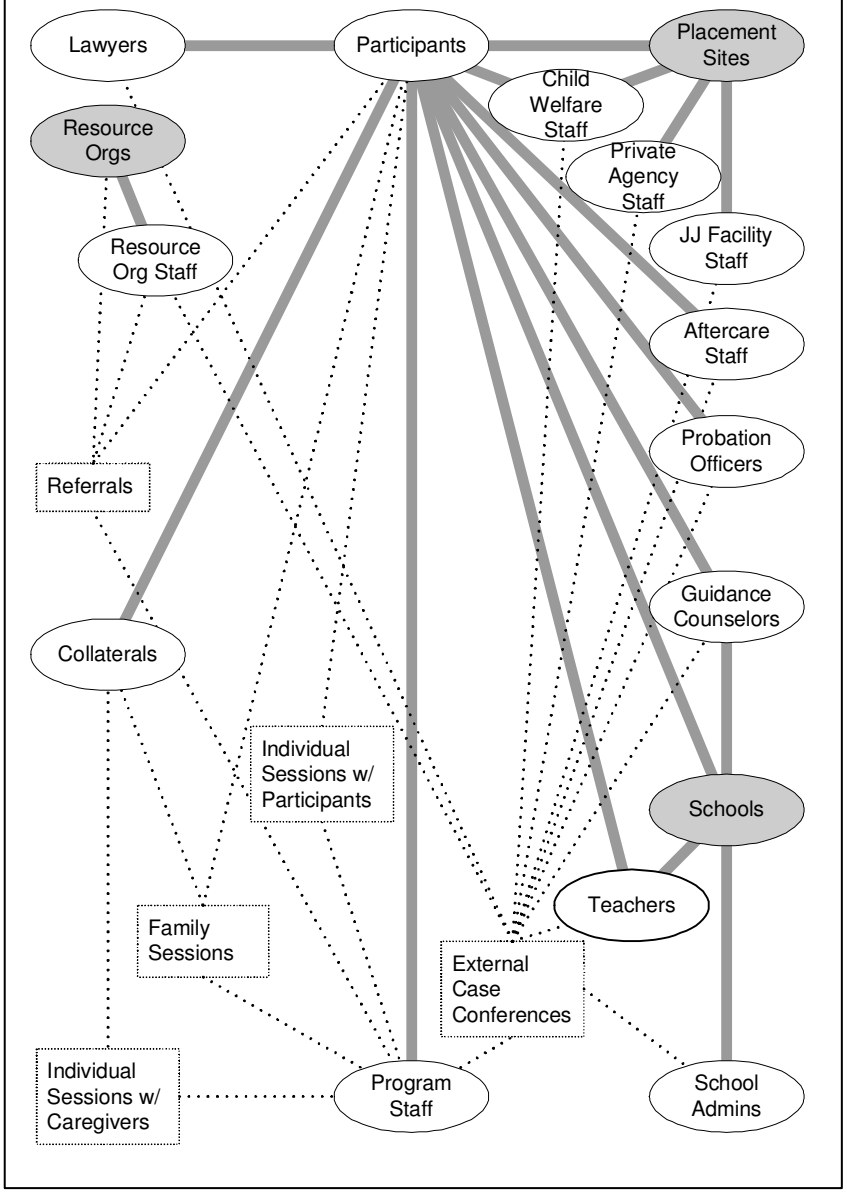

The ecosystems data model

The structure of the programmatic ecomap suggests the following semantic conventions: "A program relates to many parties, some of which are individuals and others organizations. Individuals can be further categorized according to how the program relates to them, and the same individual can fall into more than one category at any given time or over time. The same is true of organizations. One of these ways of relating to the program is for an individual to be a client. In addition to relating to the program in various ways, parties have formal relationships with each other of

various kinds and durations that are important to the program. In serving a client, the program organizes events that involve the client and other parties; the program tracks these and other events that are relevant to the client's well-being."

These statements lead to a data model that has, at its core, a generalization hierarchy of three levels. At the top is the supertype PARTY, which has two exclusive subtypes PERSON and ORGANIZATION, each of which may have multiple inclusive subtypes beneath it. One of the subtypes of PERSON will, in every case, be CLIENT. The potentially complex categorization of particular parties is handled by a many-to-many relationship between PARTY and PARTY TYPE. The structure entity RELATIONSHIP links PARTY to itself. And PARTY also has a many-to-many relationship with EVENT. The entity-relationship diagram is shown in Figure 3.

This data model has three noteworthy features:

1. The generalization hierarchy subsumes both INDIVIDUAL and ORGANIZATION under the supertype PARTY, and has a third level that facilitates the storage of data elements unique to particular subtypes of these.
2. Because the structure entity RELATIONSHIP is attached to the top level of the hierarchy, it supports the representation of all kinds of relationships among people and organizations.
3. The relationship between PARTY and EVENT is many-to-many, thereby tracking every interaction of every party within the programmatic ecosystem.

These features are consistent with a school of thought that advocates generic data models, models based on more abstract concepts than are typically used in systems design. Major data modeling resources—though silent on the human services as a specific industry—suggest similar patterns for representing enterprises in general. Hay, for example, begins his discussion of the enterprise with the point that PERSON and ORGANIZATION are reasonably modeled as subtypes of PARTY (1996). This kind of approach has been employed successfully in other settings where there is a strong need to represent interconnectedness. Contact management and customer relationship management packages, for example, are often designed in this way. Similarly, James McMillan has pointed out the value for courts of having automation systems accurately represent all of the kinds of people and entities, and the interrelationships among them, that appear in legal matters (Steelman, Goerdt, & McMillan, 2000).

This generic approach to data modeling, its advocates believe, tends to produce models with

FIGURE 3. Ecosystems Data Model

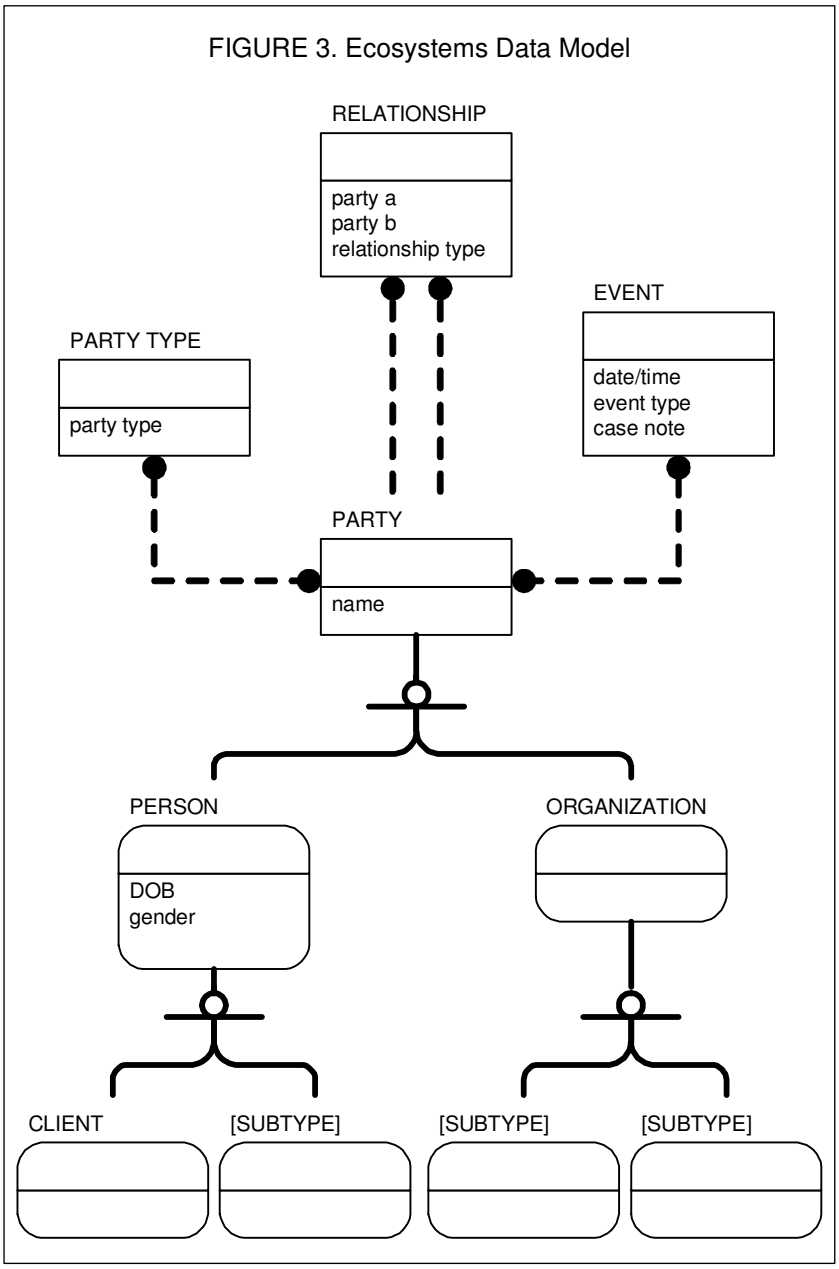

high degrees of *flexibility* and *integration*. Flexibility, in this instance, means the model's ability to deal with changes in the business or regulatory environment, while integration means consistency with data across the organization (Moody & Shanks, 2003). These qualities have long been a concern for data modeling theorists. Fleming and von Halle (1989) assert that modelers should intentionally consider not only present requirements but also how the data model could support future requirements that have yet to be formulated. Reingruber and Gregory (1994) consider a basic dimension of data model quality to be how well a model can be shared by and integrated into the work of an entire enterprise, above and beyond meeting the identified needs of a limited work area. Flexibility and integration are the qualities that make it possible for a data model to meet the common needs of ostensibly dissimilar enterprises; to meet diverse needs across the same enterprise; and to accommodate changing needs over time.

The original purpose of client-tracking software was to store case records and thereby support front-line operations and clinical supervision. Other needs, though, have emerged: providing data for program evaluators, facilitating information-sharing among program staff, and delivering indicators for performance measurement systems. The ecosystems data model offers a number of advantages in these areas.

Implications for knowledge management

The term "knowledge management" (KM) has been applied to many dissimilar endeavors, and various disparate definitions of "knowledge" have been offered, some of which are difficult to distinguish from "information" or even "data." This ambiguity has led to the charge that KM is a management consultancy fad aimed at repackaging traditional information management activities (Wilson, 2002). Though there may be a grain of truth in this criticism, there is no doubt that organizations are increasingly concerned with building strategies for effectively sharing the information garnered by individual staff members, and that this concern has led to interesting new work in software design. For the human service sector, such work has the potential to improve service delivery by bringing together relevant information that is scattered across an agency's current and former staff and even beyond the agency (Schoech, Fitch, MacFadden, & Schkade, 2002).

There are many situations in which it is useful for human service providers to understand the interconnectedness and patterns of interaction among people and organizations over time. For example, let us suppose that in Adolescent Portable Therapy, a therapist holds family therapy sessions with a participant, his parents, and his younger sister and brother; two years later, the younger sister is herself a participant and has a different therapist, and a year after that, the younger brother is also a participant with yet a different therapist. Or suppose that seven different therapists, over the course of a few years, each refers a different participant to the same resource organization, and many of those seven participants find themselves in a similar pattern of conflict with the organization's staff. In situations of this kind, the individual therapists' fragments of experience could, if brought together, coalesce in a more holistic understanding of the situation. This in turn could lead to more effective interventions.

In the natural world these fragments may come together in informal staff discussions, but that relies on chance. Formal procedures can be instituted to try to bring information together (e.g. check all case files for people of the same surname as the new participant; periodically review the set of referrals to each organization) but these will be time-consuming and therefore expensive unless an information system has been specifically designed to support them. Herein lies the value, for knowledge management, of the ecosystems data model.

In the example of the three sibling participants, because the younger sister and brother attended earlier family therapy sessions with the older brother, each of them will already have a record in the system (in the PARTY and PERSON tables) when they return as clients. And their records will already be linked to the specific sessions that each attended. This is a key advantage of the many-to-many relationship between PARTY and EVENT: it makes it possible to look specifically at the younger siblings' earlier involvement, even though they happened to first interact with the program in the context of their older brother's case. In traditional data models where EVENT is dependent on CLIENT, the siblings' earlier involvement would be hidden, buried in the case notes under the elder brother's record.

In the example of the resource organization, all referrals and external case conferences involving the organization would be recorded in the EVENT table, which would thereby

link the PARTY records of the client, the organization, and the staff member.

Because the data model captures all of these linkages, it is easy to create interfaces that display interconnectedness to the users. Interfaces from the Vera Institute's client-tracking software, as used in the APT program, are shown below.

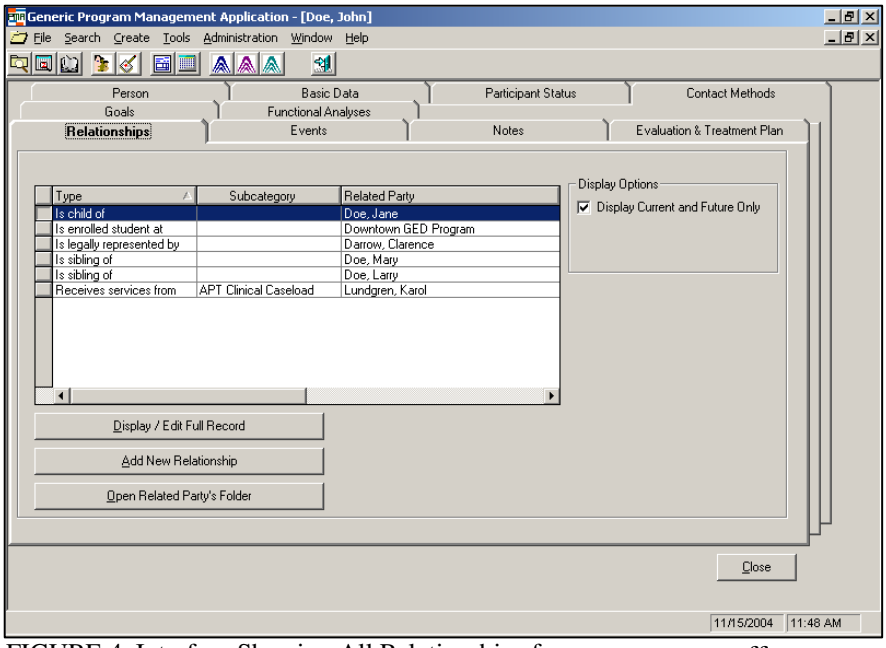

FIGURE 4. Interface Showing All Relationships for a Client

Each party—whether person or organization—has a folder. The folder includes a "Relationships" tab which shows all of the party's relationships with other parties. In this instance, the client's Relationships tab shows his APT therapist, his lawyer, his mother, his brother and sister, and the school in which he is enrolled (Figure 4). The user can begin in the client's folder and then choose to move to the folder of any of these other parties. For example, the user could move to the folder of the Downtown GED Program. There, the user would see that organization's own Relationships tab, displaying links to all of the clients who have been enrolled there, as well as to the school's teachers and other staff members.

The "Events" tab shows every event the party has ever been involved in. In the example of the new therapist whose new client had attended her brother's family therapy sessions two years ago, the therapist will see those events the first time he opens the client's folder. Likewise, a manager wishing to review the program's experience in referring clients to a resource organization over the course of several years can open the organization's "Events" tab (Figure 5) and, from there, open any of the phone calls and meetings involving that organization.

In these ways, the ecosystems data model leads to a client-tracking system that helps staff share information, and that preserves institutional memory about how people and organizations have been involved with the program across different cases. This knowledge management capability is particularly appealing because it is free. Stenmark and Lindgren (2004) note that many KM efforts falter because they require extra staff resources for maintaining the knowledge base; they suggest, as a design principle, that KM systems take advantage of spin-off from staff members' everyday work. The KM capabilities described above illustrate this principle, as they

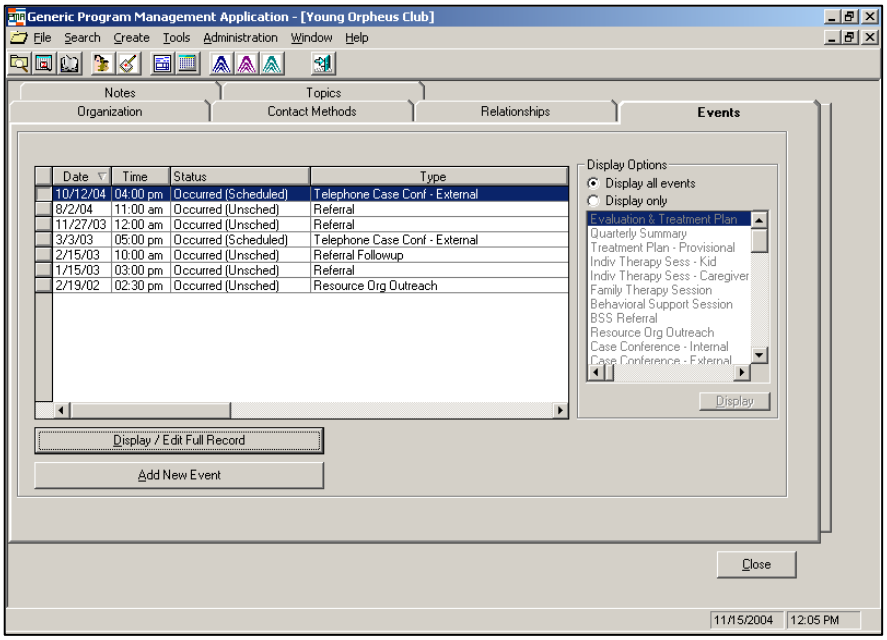

FIGURE 5. Interface Showing All Events for an Organization

derive from the data model of software used to support daily operations.

Implications for performance measurement

Many intervention models influenced by the ecosystems perspective emphasize engagement with family members and others in the client's environment. In designing performance measurement systems for such programs, it would seem reasonable to include indicators of the rate at which staff maintain contact with this constellation of people. For example, a program might set the goal of identifying three key people in each client's life, and having at least two contacts per month with each person. It should then be possible to measure the percentage of clients whose constellations received the target number of contacts.

In order to deliver this indicator, however, a client-tracking system must have a data model that explicitly distinguishes the person with whom each contact was made. It is not enough to have a Family/Collateral Phone Call record under the client's folder, and to then write out, in the case note, the name of the person called. That arrangement only permits the system to report how many family contacts were made overall regarding the client; it does not allow contacts to be grouped and counted by person to see whether there has been a steady level of contact with multiple parties across the client's constellation.

Because the ecosystems data model has a many-to-many relationship between PARTY and EVENT, it easily supports calculation of precise statistics on the rate of contact with each non-client constituent. The data model thus makes it practical to develop new kinds of metrics on a program's success in engaging with the client's ecosystem.

Conclusion

As human service databases evolve toward more accurate and complete ways of representing interconnectedness, it would be helpful if there were a more active conversation on best practices in this area. As a way of representing interconnectedness based on social work theory, the ecosystems model offers an approach that should be accessible to human service practitioners as well as information technology professionals. The advantages of the ecosystems data model demonstrated in this paper invite further discussion of the knowledge management and performance measurement capabilities that client-tracking software can provide.

References

Bilson, A., & Ross, S. (1999). Social Work Management and Practice: Systems Principles (2nd ed.). Philadelphia, PA: Jessica Kingsley Publishers.

Coursen, D., & Ferns, B. (2004). Modeling Participant Flows in Human Service Programs. Journal of Technology in Human Services, 22(4).

Fleming, C. C., & von Halle, B. (1989). Handbook of Relational Database Design. Reading, MA: Addison-Wesley.

Hartman, A. (1978). Diagrammatic assessment of family relationships. Social Casework, 59(8), 465-476.

Hay, D. C. (1996). Data Model Patterns: Conventions of Thought. New York: Dorset House.

Lohmann, R. A., & Lohmann, N. (2002). Social Administration. New York: Columbia University Press.

Mattaini, M. A. (1993). More Than a Thousand Words: Graphics for Clinical Practice. Washington, DC: NASW Press.

Meyer, C. H., & Mattaini, M. A. (1998). The Ecosystems Perspective: Implications for Practice. In M. A. Mattaini, C. T. Lowery & C. H. Meyer (Eds.), The Foundations of Social Work Practice (pp. 3-19). Washington, DC: NASW Press.

Mills-Groninger, T. (2000, November 25). Case Management data model - part 1. Message posted to the Human Service and Information Technology Association's HUSITA electronic mailing list, archived at http://listserv.uta.edu/archives/husita.html .

Mills-Groninger, T. (2003, May 1). Special Report: Change Management in Case Management [Electronic version]. TheNonProfitTimes, 19. Retrieved August 23, 2004, from http://www.nptimes.com/May2003/sr2001.html .

Moody, D. L., & Shanks, G. G. (2003). Improving the quality of data models: empirical validation of a quality management framework. Information Systems, 28, 619-650.

Reingruber, M. C., & Gregory, W. W. (1994). The Data Modeling Handbook: A Best-Practice Approach to Building Quality Data Models. New York: John Wiley & Sons.

Rosenthal, H. (2003). Human Services Dictionary. New York: Brunner-Routledge.

Schoech, D., Fitch, D., MacFadden, R., & Schkade, L. L. (2002). From Data to Intelligence : Introducing the Intelligent Organization. Administration in Social Work, 26(1), 21.

Silverston, L. (2001a). The Data Model Resource Book, Vol. 1: A Library of Universal Data Models for All Enterprises. New York: John Wiley & Sons.

Silverston, L. (2001b). The Data Model Resource Book, Vol. 2: A Library of Universal Data Models for Specific Industries. New York: John Wiley & Sons.

Steelman, D. C., Goerdt, J. A., & McMillan, J. E. (2000). Caseflow Management: The Heart of Court Management in the New Millennium. Williamsburg, VA: National Center for State Courts.

Stenmark, D., & Lindgren, R. (2004). Integrating knowledge management systems with everyday work: design principles leveraging user practice. 37th Hawaii International Conference on System Sciences.

Wilson, T. D. (2002). The nonsense of 'knowledge management'. Information Research, 8(1). Retrieved January 6, 2004, from http://informationr.net/ir/8-1/paper144.html

This is an electronic version of an article published in the Journal of Technology in Human Services 24:1 (2006). JTHS is available online at:
http://www.informaworld.com/smpp/content~db=all~content=a902286723