

Simplifying Context-Aware Agent Coordination Using Context-Sensitive Data Structures

Jamie Payton¹, Gruia-Catalin Roman¹, and Christine Julien²

¹ Department of Computer Science and Engineering
Washington University in St. Louis
{payton, roman}@wustl.edu

² Department of Electrical and Computer Engineering
The University of Texas at Austin
c.julien@mail.utexas.edu

Abstract. Context-aware computing, an emerging paradigm in which applications sense and adapt their behavior to changes in their operational environment, is key to developing dependable agent-based software systems for use in the often unpredictable settings of ad hoc networks. However, designing an application agent which interacts with other agents to gather, maintain, and adapt to context can be a difficult undertaking in an open and continuously changing environment, even for a seasoned programmer. Our goal is to simplify the programming task by hiding the details of agent coordination from the programmer, allowing one to quickly and reliably produce a context-aware application agent for use in large-scale ad hoc networks. With this goal in mind, we introduce a novel abstraction called *context-sensitive data structures* (CSDS). The programmer interacts with the CSDS through a familiar programming interface, without direct knowledge of the context gathering and maintenance tasks that occur behind the scenes. In this paper, we define a model of context-sensitive data structures, and we identify key requirements and issues associated with building an infrastructure to support the development of context-sensitive data structures.

1 Introduction

In recent years, communication technology has begun to reflect the dynamic nature of society, with devices becoming increasingly portable and untethered. The widespread use of mobile devices brings about an increased demand for software designed with mobility in mind. In fact, we can expect the number of software systems designed for use in ad hoc networks to experience rapid growth. In such networks, connections are formed opportunistically between devices within wireless communication range. Applications for this environment are likely to come into routine usage in situations such as disaster recovery in which rescue workers must find and treat victims, construction supervision in which a foreman gathers information around a site to gauge progress, etc. These and other applications for ad hoc networks are often composed from several application agents that

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE Simplifying Context-Aware Agent Coordination Using Context-Sensitive Data Structures				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Washington University, St. Louis, Department of Computer Science and Engineering, One Brookings Drive, St. Louis, MO, 63130-4899				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Context-aware computing, an emerging paradigm in which applications sense and adapt their behavior to changes in their operational environment, is key to developing dependable agent-based software systems for use in the often unpredictable settings of ad hoc networks. However, designing an application agent which interacts with other agents to gather, maintain, and adapt to context can be a difficult undertaking in an open and continuously changing environment, even for a seasoned programmer. Our goal is to simplify the programming task by hiding the details of agent coordination from the programmer, allowing one to quickly and reliably produce a context-aware application agent for use in large-scale ad hoc networks. With this goal in mind, we introduce a novel abstraction called context-sensitive data structures (CSDS). The programmer interacts with the CSDS through a familiar programming interface, without direct knowledge of the context gathering and maintenance tasks that occur behind the scenes. In this paper, we define a model of context-sensitive data structures, and we identify key requirements and issues associated with building an infrastructure to support the development of context-sensitive data structures.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

must operate in open and highly dynamic environments characterized by unpredictable and transient interactions, making it difficult for the agent programmer to produce reliable and dependable software.

Context-aware computing has been advocated as a solution for managing the programming complexity associated with such development efforts. Context-awareness refers to the ability of a software system to adapt its behavior in response to environmental changes. Typical examples of context-aware systems include location-aware offices (e.g., Active Badge [?] and PARCTab [?]), context-sensitive tour guides (e.g., Cyberguide [?] and GUIDE [?]), and context-aware note tools (e.g., FieldNote [?]). Constructing such systems is a daunting task, requiring the developer to consider the interaction between the system and a number of possibly heterogeneous sensors to gather and deliver context information.

Several frameworks and infrastructures have been devised to promote efficient, reliable development of context-aware applications by masking the complexity of interacting with heterogeneous sensors, e.g., the Context Toolkit [?] and the Context Fabric [?]. While these support systems simplify interactions with sensors, the programmer must still know the source of data to access and operate on it. In an ad hoc network, the open and dynamic nature of the environment makes it unreasonable to assume advance knowledge of the identities of data sources; application agents for use in such scenarios require a highly decoupled method of data access. Mobile agent middleware systems have been introduced that provide decoupled communication in ad hoc networks, including LIME [?], Limone [?], and EgoSpaces [?]. Many of these systems, however, are tied to the tuple space data abstraction, requiring the agent programmer to use tuple spaces as the core data abstraction in order to reap the benefits of simplified, decoupled agent communication and coordination offered by the middleware.

To provide a more general and flexible method of simplified and decoupled data access for agents operating in an ad hoc network, we propose the concept of context-sensitive data structures as the basis for a new programming methodology. A context-sensitive data structure (CSDS) is determined by and provides an agent with access to data available in the context; it is encapsulated as an abstract data type (ADT), which is represented by a class in a programming language such as Java or C++. Like all classes, it provides the agent programmer with an application programming interface (API) to access and manipulate data. The collection of data items operated on by an instantiation of such a class changes with the content of the ad hoc network. The distributed data items are accessed using the API of the local class instantiation.

The resulting design methodology provides the designer with the flexibility to use familiar and proven programming tools, i.e., ADTs, for context-aware application development. The programming tasks associated with gathering, maintaining, and adapting to context are simplified for the developer, which allows the focus to be shifted to satisfying domain-specific requirements. While implementations of context-sensitive data structures may be useful to the burgeoning

community of context-aware application developers, requiring programmers to construct an entire library of these data structures from scratch is impractical. Our goal is to provide a general model and infrastructure to support the gradual development of a library of context-sensitive data structures, which can, in turn, be used to support the context-sensitive data structures programming methodology. In this paper, we lay the conceptual foundation required to support the methodology by defining the context-sensitive data structures model and by exploring the needs of CSDS developers.

The remainder of this paper is organized as follows. Section 2 summarizes the computational model and the notion of context assumed in this paper. A motivating example of a CSDS and its use in developing a context-aware application agent is given in Section 3. Section 4 addresses the key elements required in an infrastructure for supporting the development of context-sensitive versions of traditional data structures and discusses issues with developing protocols for inclusion in the infrastructure. A brief comparison with related work is given in Section 5. The merits of the CSDS approach are discussed in Section 6, and conclusions appear in Section 7.

2 Context-Sensitive Data Structures Explained

As we embark on an exploration of the context-sensitive data structures model, we should be more specific about the environment in which an application operates. We consider agents as the main computational entities of a system, as well as providers and users of data items. To put it simply, agents are pieces of code that make up an application; each agent has its own thread of execution. Pieces of data generated by an agent are context items; context items have a general representation and can capture a wide range of information that may be important to an application, e.g., sensor readings, location information, etc. Runtime support for agents is provided by devices, which simply serve as containers for agents; we often refer to such devices as hosts. Hosts perform no application execution and do not provide or use data. An agent may be mobile, i.e., it can migrate between connected hosts. Hosts are connected when they are within wireless communication range, and agents are connected when they reside on the same host or on connected hosts. This definition of agent connectivity is important in our definition of an agent’s context. In the systems that we consider, each agent is associated with an individual context; an agent’s *maximal context* consists of context items provided by connected agents. An agent can be context-aware, meaning that it can sense changes in its context and adapt its behavior in response.

Context-sensitive data structures are an appropriate abstraction for accessing and operating on the data available in a context-aware agent’s context, and their use can greatly simplify agent interaction and collaboration. In the CSDS methodology, agent interaction occurs through the provision and acquisition of available context items. Rather than requiring an agent to collect context items from a number of agents spanning an ad hoc network or provide a context

item to a particular agent, the CSDS methodology supports a more decoupled style of agent interaction. With our approach, an agent declares and instantiates a context-sensitive data structure. The content of the context-sensitive data structure is determined by the agent's associated context. Given that an agent's maximal context can span an entire large scale multi-agent system, the context can grow quite large and unmanageable. As pointed out in [?], one way to allow a context-aware agent to economically manage its context is to supply the agent with a more limited context that is tailored to its needs. In the context-sensitive data structures programming methodology, an agent can provide a tailored context definition in the instantiation of a CSDS, essentially specifying the desired context for the data structure to operate over.

When a context-sensitive data structure is used by an agent, the task of managing access to the data elements that are spread across the ad hoc network while maintaining a specific data organization defined by the structure is hidden from the application programmer. Access to the data elements of the CSDS is gained only through operations defined on its ADT. Operations performed on the CSDS can effect a change in the context of others, as can the movement of an agent that may cause it to join or leave another agent's context. As these changes in the state of the environment occur, the content of the context-sensitive data structure is changed appropriately in response. A developer using a context-sensitive data structure can operate on the dynamically changing set of data elements that are distributed throughout the context as if the data were stored in a local, persistent data structure. The management of the data elements within the CSDS is automatically handled in the face of changes without intervention by the application programmer, since the data structure is essentially a reflection of context.

In the remainder of this paper, we investigate the software engineering potential for context-sensitive data structures. First, we offer a concrete example of a context-aware application that can benefit from the use of a particular context-sensitive data structure, the priority queue. We then explore protocols we must provide in the infrastructure to support the development of context-sensitive data structures for use in such applications. In doing so, we seek to demonstrate the feasibility of applying the context-sensitive data structures concept and associated design methodology.

3 Programming with Context-Sensitive Data Structures

The impetus behind the introduction of the context-sensitive data structure design methodology is to reduce development costs in terms of effort and errors, and to make context-aware application development accessible even to novice programmers. Context-sensitive data structures provide a decoupled method of accessing and operating on data in the ad hoc network, one that is simple and natural to the programmer, using the same interface as in static settings. Moreover, the dynamically changing content is managed transparently, reducing the complexity of the environment, and, in turn, the potential for programming er-

rors incurred by interacting with agents in a large-scale and highly dynamic ad hoc network. Finally, the context-sensitive data structures design methodology offers a flexible approach to software development for ad hoc networks, and gives the agent programmer the freedom to choose the data abstraction that is best suited for task at hand.

In the remainder of this section, we use an example to demonstrate the utility of context-sensitive data structures and the associated design methodology. The example highlights how the CSDS methodology can be applied and how a particular instance of a context-sensitive data structure can be used in an application scenario.

We consider a disaster recovery scenario in which triage is employed to treat the wounded. Victims are quickly examined to evaluate the seriousness of their injuries and are tagged with devices that emit (via wireless radio or infrared) information about the assigned injury classification, ranging from injuries that need immediate attention to those for which treatment can be postponed. Rescue teams are assigned areas in which they must arrange transport for the most severely injured first and provide as much on-site treatment as possible for these victims until transport is available. The rescue team members use PDAs with wireless communication capabilities to coordinate activities and to obtain and display the status of victims and volunteers. A volunteer is selected by the rescue team member to treat the most seriously wounded victim until transport arrives. A volunteer's assignment may change as the status of injured victims within the context changes. After a rescue crew member arranges on-site treatment for a victim, he must arrange for the victim's transport to a hospital. As victims are transported, they are removed from the context of the application. As new victims are discovered and their injuries evaluated, they are added to the context. Figure ?? illustrates this application.

Fig. 1. Disaster Recovery Scenario. The disaster site lies within the large oval. A rescue crew member (the encircled cross) uses a PDA that runs an application to assign to the most seriously wounded victims in the designated area (the dashed box) on-site treatment and ambulance transport to a nearby hospital. Victims are shown as circles, with seriousness of injury reflected by darker shading.

Building the application described from scratch can be a significant undertaking. The programmer must include functions to sense the set of neighboring hosts, to send messages to agents on reachable hosts, and to issue queries to obtain data. Given the intended pattern of data access in this particular application, the use of a priority queue data structure seems like the most intuitive way to represent data. Therefore, query responses must be processed and placed into a traditional, static priority queue. Each and every time an operation is requested, the hosts in the network must be queried to ensure operation over a set of data most closely reflecting the current state of the context, and the

replies must be processed to ensure an appropriate presentation of the resulting data. In this particular example, such processing requires proper insertion of the resulting data elements in the traditional priority queue data structure. The remainder of this section illustrates a CSDS approach to implementing the disaster recovery application and demonstrates how context-sensitive ADTs can be used by application programmers to reduce the complexity of context maintenance duties associated with context-aware programming.

A context-sensitive data structure is captured as a class implemented in a programming language. An agent's instantiation of a CSDS, at any instant in time, encapsulates the data items available in the agent's context. To be more specific, the content of the context-sensitive data structure is defined by a context specification provided by the agent programmer. The context specification gives rules on what kinds of context information should be provided to the agent in order to restrict the agent's context to a manageable size. The CSDS places a logical ordering over the data items within the restricted context to mimic the organization of its traditional data structure counterpart, even as the data items are constantly changing due to changes in the context. Access to the changing collection of data items is gained only through the CSDS API, which includes data access and manipulation operations similar to those provided for traditional data structures. When using the CSDS programming methodology, the amount of data processed by the application agent is reduced, explicit data maintenance by the application programmer is removed, and application development is simplified.

A simple application agent for rescue team support could be constructed around the notion of a context-sensitive data structure. While several different context-sensitive data structures could be employed, based on the pattern of data access recognizable in the disaster recovery scenario, we use a context-sensitive priority queue as a motivating example.

In the disaster recovery example that we consider, the agent provides a context specification with the context-sensitive priority queue instantiation that restricts the context items eligible for inclusion in the priority queue to those that lie within a limited area of the disaster site, e.g., a one block radius. The context definition restricts which elements of the entire operational environment will be included as items in the priority queue. Data elements of the priority queue used in the triage application are pieces of injury information emitted from victims' triage tags. Given the above context specification, the content of the priority queue for a crew member's application is the injury information data elements that are located within her assigned area. The data associated with the context-sensitive priority queue reflects an ordering over the injured within that area such that the most seriously injured victim is at the head of the queue. The context, and hence the content of the context-sensitive priority queue, is updated independently of the application's operation on the queue. No context management gathering or maintenance is explicitly performed by the agent that instantiates a context-sensitive data structure; in this example, the burden of

context maintenance duties falls on the context-sensitive ADT implementation of the priority queue.

In this example, our context-sensitive priority queue requires two operations: `getFirst()` and `removeFirst()`. In our application, `getFirst()` is used to access an injury description for the victim in the context with the most severe injury. The injury description includes a unique injury identifier, the injury priority, and the geographical location of the injured person. The `removeFirst()` operation is used to access the injury description of the victim with the highest priority injury and to remove the injury description from consideration of all medics that schedule transport. The victim's information is still made available to ambulance drivers. This type of removal can be accomplished simply by changing the context item containing the injury status of a victim such that it is no longer included in priority queues of medics, but is still included in the context-sensitive data structures used by ambulance drivers.

It may seem that we have omitted operations needed to populate a priority queue. While explicit data insertion operations may be needed in other applications, none are needed for this scenario. Data elements become available as a result of the introduction of devices that emit injury information, and are included in a rescue crew member's application as a result of context-maintenance performed to uphold the provided context definition. (It is important to note that while the application presented here requires only implicit insertion of data items by the infrastructure, the general context-sensitive data structures model is not limited to this type of insertion. A discussion of issues related to supporting insertion operations is presented in Section 4.)

In the disaster recovery application, the context-sensitive priority queue consists of the victims in the context ordered by injury priority. The rescue crew member uses the application to get the head of the priority queue, dispatching a volunteer to tend to the victim until transport can be arranged. Because we consider that crew members may be assigned overlapping contexts and that the transport vehicles available to one crew member may not be available to another, the injury description obtained to dispatch treatment should still be made available. For this reason, the dispatch function of the application is implemented using the `getFirst()` operation previously described. Once treatment has been dispatched to the most severely wounded victim, the crew member uses the application on his PDA to determine if any transportation resources are available. If so, the application assigns the available transportation resources to the most severely injured victim in the context. Because the victim has been assigned on-site treatment and scheduled for evacuation, the victim should be removed from consideration by the rescue crew teams. Therefore, transport scheduling in the application should be implemented using the `removeFirst()` operation.

Figure ?? shows sample code for a straightforward Java-based implementation of the disaster recovery application using a context-sensitive priority queue. This version of the agent simply defines a context, instantiates the context-sensitive priority queue, and performs processing on the priority queue using the operations made available by the API, e.g., `getFirst()` and `removeFirst()`.

```

public class DisasterRecoveryAgent extends Thread {

    public DisasterRecoveryAgent()
        Context context = one block radius
        PriorityQueue pq = new PriorityQueue(context);

    public void run()
        while(victimsUntreated())
            if(volunteersAvailable())
                TreatmentThread treat = new TreatmentThread(pq);
                treat.start();

            if(transportAvailable())
                TransportThread transport = new TransportThread(pq);
                transport.start();

    class TreatmentThread extends Thread
        the start method calls the run method...
        public void run()
            dispatch(getVolunteer(), (pq.getFirst()).id);

    class TransportThread extends Thread
        the start method calls the run method...
        public void run()
            assign(getTransport(), (pq.removeFirst()).id);

}

```

Fig. 2. A CSDS Approach to the Disaster Recovery Application

The data structure does not have to be explicitly reconfigured by the application each time a victim is transported. Instead, an untreated victim in the context with the highest injury priority can be identified simply by using the `getFirst()` operation.

This example is suggestive of the programming productivity gains one could achieve with context-sensitive data structures. A variety of context-aware applications can benefit from the use of a number of other context-sensitive data structures as well. Providing a library of context-sensitive data structures, however, requires a demanding and time-consuming process. In the next section, we explore what is needed to support the implementation of context-sensitive data structures such as the priority queue used in the disaster recovery application.

4 CSDS Infrastructure Support

We envision the gradual development of a library of context-sensitive data structures for use by context-aware application programmers. In most cases, the application programmer should not have to implement the context-sensitive data

structure; she should simply choose among the available CSDS implementations. The application programmer is expected to use the API of the selected CSDS to interact with data as if it were local. Since many data structures share common operations, we envision providing an infrastructure that supports the development of context-sensitive lists, trees, stacks, queues, and other data structures.

At the heart of the CSDS model is the perception that we are populating a locally accessible structure with data items distributed throughout the associated agent’s context, keeping the items in the local view consistent with the context as the environment changes. In reality, we are building a structure on top of the ad hoc network—a structure which mimics the organization imposed by a particular data structure. This overlay structure is used to support operations issued on a CSDS. As such, the structure must adapt accordingly in response to context changes.

To deliver a CSDS support infrastructure, we must explore what is required to build and maintain an overlay structure over the ad hoc network. To begin, we consider that the environment in which agents operate is open and dynamic. As the number of hosts that join the network grows, the number of context items available to an application agent significantly increases. Building an overlay structure to support the operation of an agent’s CSDS over a large body of context items requires a substantial amount of processing. To aid in the development of efficient context-sensitive data structures, the infrastructure must contain protocols for limiting the scope of the context to include only those items that suit an agent’s particular needs. The tailored context delivered as a result will be used by other protocols required in our infrastructure: those for supporting the implementation of particular operations on a CSDS. In the remainder of this section, we explore design issues associated with context scoping protocols and examine the effects of various CSDS operations on the development of protocols for building and maintaining overlay structures for ad hoc networks.

Before we begin this exploration, we remind the reader that our ultimate goal is to provide a general support infrastructure containing a set of protocols that can be used to aid others in the development of an entire library of context-sensitive data structures to be used by agent developers following the CSDS development methodology. As such, the discussion of data structures and their operations is generalized for dynamic set data structures. Inspiration for the generalized descriptions of data structure population, access, and removal operations is taken from [?].

4.1 Protocols for Tailored Contexts

There are several viable approaches to limiting the reach of an agent’s context. We utilize a policy-based approach similar to that in the network abstractions protocol [?] in which a context-scoping policy is used to determine an agent’s context. In our approach, a context-scoping policy is associated with a particular CSDS. The policy is used to govern which context items in the ad hoc network are eligible for inclusion in the CSDS. An application programmer can specify

the context associated with a CSDS by providing a context-scoping policy as part of the data structure’s instantiation.

Each policy is defined as a set of constraints on properties of the ad hoc network. Constraints on properties of hosts (e.g., battery life), of communication links (e.g., bandwidth), of agents (e.g., access rights), and of data (e.g., type) may be used to define a context specification policy. We favor policy specifications that use constraints on such properties because they offer generality and flexibility and allow developers to reason at a higher level of abstraction about the entities within the ad hoc network and the way they contribute to defining the content of the CSDS.

The context-scoping protocol uses the scoping policy supplied by the application programmer to present a subset of the items in the ad hoc network as the content of the CSDS. In doing so, the protocol builds a context structure over the ad hoc network, which can then be used by other protocols that support the execution of data structure operations. Certain scenarios call for different ways of using the context structure. When the environment is highly dynamic and data structure operations are issued over the context infrequently, the context structure is built on-demand each time that an operation on the associated data structure is issued. In situations where the environment is relatively stable and operations over the context are frequently performed, the context structure is maintained as the environment changes.

As a final note, implementations that supply agents with tailored contexts are implemented in a distributed fashion. Agents do not require global knowledge of the environment to participate in the computation of and to interact with their tailored context.

4.2 Data Structure Population Protocols

Typically, the insert operation described below is used to populate traditional data structures:

- **insert(*X*)**: places the data element *X* in the data structure according to its organizational policy.

Rather than allow programmers to insert data directly into a CSDS, the infrastructure performs data maintenance on behalf of the application. Thus, the **insert(*X*)** operation is not directly provided to programmers for a particular data structure in the context-sensitive data structures methodology. Instead, there are two ways to include data items in an application’s CSDS: indirectly through context-specification and data element ordering, or directly by injecting data as a context item into the environment.

First, indirect insertion is used to populate a data structure. An application-provided context-scoping policy is supplied to the infrastructure in the instantiation of a CSDS. A protocol in the infrastructure for providing tailored contexts selects the data elements to be contained within the CSDS, and a separate protocol creates an overlay structure to mimic a local organization of those elements

according to properties of the data structure. Like the context-scoping protocol, it may be practical in some situations to build the overlay structure on-demand when an operation is issued, or it may be more effective to maintain the overlay structure in the presence of changes. The overlay structure protocols often utilize several other protocols for data collection and aggregation. For instance, the implementation of the context-sensitive priority queue uses a protocol that sorts the items in the scoped context and returns the greatest element.

Second, direct insertion is performed by using an insertion operation provided on the *infrastructure* instead of on the context-sensitive data structures. We treat each data item produced by an application as a generic piece of data that is provided to the infrastructure via `insert` in order to supply it to other agents as context. The semantics of direct data insertion varies, and a suitable option can be specified by the programmer. We identify three types of direct insertion operations: local, destination-aware, and property-aware. In *local* insertions, the data item is stored locally by the inserting agent. A *destination-aware* insertion allows the programmer to specify a desired destination for a data item in terms of a particular agent or host. The inserted context item will eventually reside at the specified destination through the use of auto-migration. With auto-migration, the data item is delivered immediately when the destination is available. If the destination is unavailable (e.g., because of network partitioning), the data item is marked for migration and stored locally until it can be delivered. The *property-aware* insertion is similar, but allows a more decoupled method of specifying a recipient. With this type of insertion, a policy restricting the set of potential destinations is provided as a parameter to the insert operation. All destinations are evaluated against the criteria. A single destination is non-deterministically chosen from the set of matching destinations and is used by the infrastructure in a destination-aware insertion.

To allow agents to protect their data, specialized versions of the direct insert operations that support access control mechanisms can be used. Access control parameters are included with an insert operation to specify how the data is made available at different levels of protection. Only authorized agents are allowed to access or delete another agent's data items. The programmer can specify which agents are authorized using a policy similar to that used for context-scoping. This form of access control can be supported in part by requiring all agents to provide credentials in the context definition used to populate the data structure. These credentials are used by the context-scoping protocols to evaluate the access control rights against the access control policies of the provider to determine if the data is included in the context.

All of these insertion styles may affect the context associated with an agent's CSDS. If an overlay structure maintained for a particular context-sensitive data structure is affected by some agent's insertion operation, the protocol for maintaining the overlay structure must sense the change in the environment and accordingly adapt the structure.

4.3 Data Access Protocols

Data access in traditional dynamic set data structures can be generalized by the set of operations described below:

- **get(X)**: searches the data structure for the item corresponding to the key X . If successful, the operation returns the corresponding element; otherwise, it returns *null*.
- **contains(X)**: searches the data structure for the item corresponding to the key X . If successful, the operation returns *true*; otherwise, it returns *false*.
- **getNext()**: returns the next data element in the data structure. If the element does not exist, *null* is returned.
- **getFirst()**: returns the data element located in the first position of the data structure. If the element does not exist, *null* is returned.
- **iterate()**: returns an iterator over the data structure.

These operations do not change the data structure in any way. Thus, protocols designed to support data access operations simply use the overlay structure built by the population protocols discussed in the previous subsection. Examining the set of operations brings up questions about the semantics provided by the protocols. In some situations, an application's requirements may be satisfied by weakly consistent results in exchange for more efficient execution of operations. In other scenarios, a strongly consistent reflection of the environment is required in the result, regardless of the expense of the distributed transactions needed for the operation's execution.

4.4 Data Removal Protocols

Manipulating the data structure by removing elements is a common task, and is typically achieved through the use of operations such as:

- **remove(X)**: returns and deletes the data element X from the data structure if it exists and adjusts the data structure if needed. If the element X does not exist in the data structure, the operation returns *null*.
- **removeNext()**: returns and deletes the next data element from the data structure if it exists and adjusts the data structure if needed. If the element does not exist, the operation returns *null*.
- **removeFirst()**: returns and deletes the data element located in the first position of the data structure if it exists, and adjusts the data structure; otherwise, the operation returns *null*.

Protocols developed to support these removal operations may have different semantics. We provide two types of removal operations: an *individual* remove and a *communal* remove. The former eliminates a data element from inclusion only for the issuing agent's particular CSDS on which the operation was called, while the latter expunges the data item from inclusion in any CSDS of any agent by removing the data item from the ad hoc network.

Individual remove operations are useful for collaborative applications that operate on overlapping contexts, such as the disaster recovery scenario presented earlier. In this application, once a rescue team member arranges treatment and transport, the victim is removed from the context-sensitive priority queue. However, volunteers and ambulance crews still require access to the injury information, and so it is not removed from the ad hoc network. To support individual removal operations, the protocol performs bookkeeping. When a remove is issued, the specified data item is marked as no longer belonging to a particular CSDS and that information is remembered by the owner as part of the data element. Choosing to use this protocol in the development of a CSDS, however, requires careful consideration, as the bookkeeping required can create a significant amount of overhead.

The more traditional communal remove operation eradicates the specified element both from the context-sensitive data structure and from the ad hoc network. This approach also requires careful consideration, since the protocol essentially deletes another agent’s data. Using the access control approach previously mentioned, however, allows agents to control how other agents access their data.

Regardless of choice between individual or communal removal semantics, if the overlay structure is maintained, the removal operations require its restructuring. The removal of a piece of data is essentially a change in context and is handled by rebuilding the overlay structure when a data structure operation is issued over the context, or by the overlay structure maintenance protocol discussed in subsection ??.

4.5 Implementation Requirements and Issues

As we explore the potential for CSDS development, we make the observation that with particular data structures, the same portion is regularly accessed. For instance, priority queues and stacks are frequently accessed at the beginning positions of the data structures, using operations such as `getFirst` and `pop`. This observation is the motivation behind the concept of on-demand partially maintained data structures. We believe that improvement upon the performance of a typical CSDS can be facilitated by relaxing the requirement that the structure be built and maintained over all the data items in the context. Instead, a CSDS is initially constructed to consist of only the first n elements, where n is a parameter given in the instantiation of the data structure. As the elements are accessed, the structure is further constructed on-demand. The structure is maintained to n elements in the presence of context changes. The parameter n is tunable and can be changed to address the application’s need or changes in the context. Our approach to partially maintained data structures has its roots in the *suspended cons* concept described in [?], which was introduced to support finite storage of infinite objects. Suspended `cons` is an extension to Lisp that allows placeholders of expressions to be stored until an operation forces its evaluation. Similarly, our partially maintained context-sensitive data structures are evaluated further only upon demand.

As a final note on supporting the development of context-aware data structures, it is imperative that we carefully evaluate the requirements and issues presented, develop protocols in response, and include them in an infrastructure. The delivered infrastructure should be flexible, allowing the CSDS programmer to use only needed components. The components should have minimal programming interfaces which are familiar and intuitive. The protocols should perform in a reasonably efficient manner. Moreover, the CSDS design methodology for context-aware application development should be put to the test through the development of applications that use context-sensitive data structures.

5 Related Work

One approach to simplifying context interaction is to treat the collection of data in the ad hoc network as a database. For instance, the TAG system [?] was developed to simplify sensor network data collection for programmers by providing a declarative interface in the form of a simple query language. The query language is based on SQL, with the main difference being that persistent queries are supported through the use of an additional clause that allows a user to specify a duration for periodic evaluation. Further work resulted in the TinyDB system [?] which employs an acquisitional query processing approach in which the sensors themselves can assist in intelligent collection of data to reduce power consumption. While these and other such approaches simplify context interaction and data collection in dynamic environments, they may not be the solution to supporting rapid development of multi-agent systems. This kind of system still requires awareness of the environment; the programmer must have some knowledge of how the data in the network is structured within the virtual “database” in order to issue meaningful queries. The application programmer is also responsible for processing the data obtained from the network using such queries and organizing it into an appropriate data structure. In addition, though the specification of a duration allows for periodically evaluated persistent queries, such an approach does not allow for a more general specification of consistency between the data residing in the network and the returned results. Finally, the database approach was intended solely for data collection, and does not lend data encapsulation as does a data structures approach.

Closer in spirit to our work are those approaches which rely on a data structure abstraction to support rapid development of context-aware applications. In fact, a simple example of the context-sensitive data structures concept can be found in models which utilize a transiently shared tuple space data structure to support coordination among applications in ad hoc networks, e.g., LIME [?], Limone [?], and others. Another example can be found in the EgoSpaces [?] model, which relies on a slightly more sophisticated abstraction to facilitate agent coordination in ad hoc networks. In this model, an agent’s tailored, asymmetric context is encapsulated in an abstraction called a view. Agent interaction occurs through the use of insertion and content-based retrieval operations on the view. The middleware implementations of these models are strongly tied to a

single data abstraction, e.g., the tuple space in LIME and the view in EgoSpaces. In contrast, in our work we seek a more general approach in order to provide programmers with the ability to use a range of context-sensitive versions of traditional data structures, such as stacks, queues, and trees, allowing the programmer to choose the most appropriate data structure which best suits the needs of the context-aware application.

Essential to making our approach viable is the ability to restrict an application agent’s context to a manageable subset of the data available in the network. As previously mentioned, one protocol that allows for declarative specification and maintenance of context tailored to an agent’s particular needs is the network abstractions protocol [?]. While the network abstractions protocol provides the ability to define a number of useful contexts, it is not possible to define all contexts that an application desires. Consider, for instance, a city employee who wants to monitor water meters distributed throughout the city. The context for his application could be defined as “all meters until a meter outside the city limits is reached”. Another application might require a context based on temporal properties. For instance, an application that uses temperature data in the surrounding area to adapt its operation may only want to act upon data readings that are relatively fresh. To our knowledge, no protocols exist to define these kinds of contexts. We would like to develop protocols that allow specification of these and other contexts and to include them in our infrastructure.

6 Discussion

The purpose of the context-sensitive data structures approach is to support rapid development of context-aware applications. However, programmer productivity is not our only concern. In the remainder of this section, we discuss how the CSDS programming methodology and the support system outlined in this paper meet quality requirements associated with the development of large-scale multi-agent software systems.

Scalability. Since we consider multi-agent systems that operate in open environments, the number of agents in the network can grow large. As such, an agent’s maximal context (which spans the entire network) can include context information provided by a considerable number of hosts. The result is that a context-aware agent may be hindered rather than helped in its performance by the use of context, since the agent becomes consumed by an inordinate amount of context collection and processing. The context-sensitive data structures approach considers the need to keep an agent’s context manageable by allowing the programmer to specify restricted contexts, and to associate a particular context with a context-sensitive data structure.

Dependability. A common approach to bolstering the dependability of multi-agent systems that operate in unpredictable environments is to employ the notion of context-awareness in agent development. Using the context-sensitive data structures approach simplifies dependable context-aware agent development by automating the context management tasks and hiding them from the

agent programmer, presenting instead a familiar interface in the form of data structures. It follows that programming errors associated with context collection tasks can be reduced when the context-sensitive data structures methodology is utilized.

Reusability. The context-sensitive data structures programming methodology advocates a style of context-aware agent system development that supports reusability. Encapsulating all context management tasks in a CSDS eliminates the need for rewriting context management code across applications—an agent programmer can instead reuse the CSDS. By allowing a CSDS to be associated with a particular context specification at instantiation, a CSDS can be reused for different kinds of contexts. Furthermore, since population and maintenance of the contents of a CSDS is determined by the context specification provided in the data structure’s instantiation, agent code can be reused for a different purpose simply by changing the CSDS’s associated context specification.

While context-sensitive data structures offer many benefits in developing context-aware multi-agent systems, it remains to be seen how effective their use can be. Much work lies ahead in the development and evaluation of efficient protocols for inclusion in the CSDS infrastructure, as well as the development of a CSDS library.

7 Conclusions

In this paper, we presented a novel abstraction called the context-sensitive data structure designed to simplify the development of context-aware applications. A context-sensitive data structure encapsulates data items distributed among a number of agents within a restricted portion of the large-scale ad hoc network and provides the programmer with access to the collection of data elements as if they were local through a well-defined API. The content of the context-sensitive data structure is fluid; as the context changes, the CSDS is reorganized to reflect the changes in the state of the environment. To support the use of context-sensitive data structures as a design methodology, we proposed providing an infrastructure that encapsulates protocols for restricting the context, for accessing data elements in the context, and for modifying data elements in the context. We envision this infrastructure as providing the CSDS developer with a set of essential tools that can be used to develop a range of context-sensitive data structures. In this paper we take a first step toward that goal by defining the context-sensitive data structures model and outlining the requirements and issues associated with developing a CSDS support infrastructure.

Acknowledgements

This research was supported in part by the Office of Naval Research under ONR MURI research contract N00014-02-1-0715. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsoring agencies.