

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCSE-2003-66

2003-10-02

A Unified Specification Framework for Spatiotemporal Communication

Gruia-Catalin Roman, Octav Chipara, Chien-Liang Fok, and Chien-Liang Fok

Traditionally, network communication entailed the delivery of messages to specific network addresses. As computers acquired multimedia capabilities, new applications such as video broadcasting dictated the need for real-time quality of service guarantees and delivery to multiple recipients. In light of this, a subtle transition took place as a subset of IP addresses evolved into a group-naming scheme and best-effort delivery became subjugated to temporal constraints. With recent developments in mobile and sensor networks new applications are being considered in which physical locations and even temporal coordinates play a role in identifying the set of desired recipients. Other applications involved... **Read complete abstract on page 2.**

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research

Recommended Citation

Roman, Gruia-Catalin; Chipara, Octav; Fok, Chien-Liang; and Fok, Chien-Liang, "A Unified Specification Framework for Spatiotemporal Communication" Report Number: WUCSE-2003-66 (2003). *All Computer Science and Engineering Research*.

https://openscholarship.wustl.edu/cse_research/1112

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

A Unified Specification Framework for Spatiotemporal Communication

Gruia-Catalin Roman, Octav Chipara, Chien-Liang Fok, and Chien-Liang Fok

Complete Abstract:

Traditionally, network communication entailed the delivery of messages to specific network addresses. As computers acquired multimedia capabilities, new applications such as video broadcasting dictated the need for real-time quality of service guarantees and delivery to multiple recipients. In light of this, a subtle transition took place as a subset of IP addresses evolved into a group-naming scheme and best-effort delivery became subjugated to temporal constraints. With recent developments in mobile and sensor networks new applications are being considered in which physical locations and even temporal coordinates play a role in identifying the set of desired recipients. Other applications involved in the delivery of spatiotemporal services are pointing to increasingly sophisticated ways in which the name, time, and space dimensions can be engaged in specifying the recipients of a given message. In this paper we explore the extent to which these and other techniques for implicit and explicit specification of the recipient list can be brought under a single unified framework. The proposed framework is shown to be expressive enough so as to offer precise specifications for existing communication mechanisms. More importantly, its analysis suggests novel forms of communication relevant to the emerging areas of spatiotemporal service provision in sensor and mobile networks.

A Unified Specification Framework for Spatiotemporal Communication

Gruia-Catalin Roman, Octav Chipara, Chien-Liang Fok,
Qingfeng Huang¹, and Chenyang Lu

Department of Computer Science and Engineering
Washington University
Campus Box 1045, One Brookings Drive
St. Louis, MO 63130-4899, USA
{roman, ochipara, liang, qingfeng, lu}@cse.wustl.edu

Abstract

Traditionally, network communication entailed the delivery of messages to specific network addresses. As computers acquired multimedia capabilities, new applications such as video broadcasting dictated the need for real-time quality of service guarantees and delivery to multiple recipients. In light of this, a subtle transition took place as a subset of IP addresses evolved into a group-naming scheme and best-effort delivery became subjugated to temporal constraints. With recent developments in mobile and sensor networks new applications are being considered in which physical locations and even temporal coordinates play a role in identifying the set of desired recipients. Other applications involved in the delivery of spatiotemporal services are pointing to increasingly sophisticated ways in which the name, time, and space dimensions can be engaged in specifying the recipients of a given message. In this paper we explore the extent to which these and other techniques for implicit and explicit specification of the recipient list can be brought under a single unified framework. The proposed framework is shown to be expressive enough so as to offer precise specifications for existing communication mechanisms. More importantly, its analysis suggests novel forms of communication relevant to the emerging areas of spatiotemporal service provision in sensor and mobile networks.

1. Introduction

Networking is a mature field with a major impact on the society as a whole. It connected the world in ways never before imagined and made an entire generation addicted to being in contact with each other and with all that is happening around the globe. Key to the notion of communication is the simple idea of transmitting a message from one point to another. Today's networking infrastructure has been the direct result of a gigantic engineering effort to achieve reliable and fast transmission of the individual message. Transport protocols are the conceptual bridge between the network fabric and the application layer that exploits it. They reflect, in part, low-level communication requirements distilled from needs that manifest themselves at the application level. Unicast captures the notion of delivering a message to a known destination. Multicast relates to the desire to provide the same information to an entire community of hosts or users. Anycast expresses the notion that (within reasonable limits) the identity of the recipient is not critical as long as someone receives the message. Because networking is, par excellence, a service oriented technology, its ability to deliver changes the level of expectation within the user community and the applications that serve it. This is visible not only in the growing demand for increased performance but also in a gradual shift of the application profile and the subsequent pressure for offering new transport protocols and paradigms.

Video and audio streaming requirements make explicit the temporal dimension of message delivery and led to the emergence of an entire new class of proto-

¹Currently at Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94040

cols. Location sensitive Internet applications designed to identify resources within proximity of the query initiator resulted in the development of protocols that have a spatial dimension. A new way of thinking about how to specify the destination of a message is emerging. The notion of a multicast address created one level of indirection in the naming of the destination nodes; only nodes that subscribed to the particular multicast receive the message. Geocast offered another style of indirect specification based on the presence within a particular location in space. Nomadic networks, in which mobile nodes on the fringe of the wired network communicate with the rest of the world via base stations, as in Mobile IP, are reached by contacting a home agent who knows how to deliver the message to the remote location.

This process of innovation shows no signs of slowing down with many new protocols being proposed in recent years. The high level of activity is fueled, in part, by a combination of societal changes and technological advances in wireless communication and sensing devices. Here are some examples that show how the desire to exploit mobile and sensor networks can lead to the definition of novel transport protocols. Let us consider the case of a soldier who is running through an unfamiliar area dotted in advance by ground sensors dispersed from a helicopter. For safety reasons, he would like to know of any possible threats lying ahead along his path. The soldier's PDA can send a scouting request to a delivery zone that moves on his path in front of him. Only the sensors that enter the delivery zone receive the scouting message, pool their currently sensed information, and send the aggregated data back to him. A recently proposed new protocol called mobicast naturally delivers the spatial and temporal locality requirements of information dissemination and gathering exhibited by this application. It does so by allowing one to specify a message delivery zone that evolves over time. This is possibly the first instance when the temporal and spatial dimensions are brought together to define the set of intended recipients. Other protocols motivated by developments in the area of sensor networks involve area-based anycast and predictable message delivery.

In a field not known to be overly concerned with formal methods, this plethora of novel protocols raises serious questions about their precise specification, about the relation among the different proposals, and about unexplored research opportunities. Any effort designed to organize this field in a principled way has the potential to contribute in a timely fashion to an important research direction. Bringing the software engineering perspective to this process is particularly significant in

freshly evolving areas such as mobile and sensor networks. In this paper, we seek to accomplish this by exploring the specification dimension of protocol definition. The goal is to create a unified framework that could serve as the formal foundation for categorizing both established and nascent protocols and for investigating opportunities for future research in this area.

We started our investigation from the observation that, as a whole, existing protocols touched upon three important specification dimensions in their approaches to defining the set of intended message recipients: the naming, time, and space domains. The immediate generalization that resulted from this way of thinking was the notion of a delivery volume in this three-dimensional space. In the simplest of terms, this relates to the notion of delivering to certain nodes present in certain places at certain times.

The remainder of the paper formalizes this concept, explores various instantiations and possible generalizations, and relates it to existing work in networking protocols. Section 2 introduces the specification notation, its formal semantics, and presents a simple example of how it is used to capture IP unicast. Section 3 shows how our specification notation is used to capture the delivery semantics of a variety of protocols. It also discusses opportunities for achieving a higher levels of generality. Section 4 concludes the paper.

2. Basic Concepts

Advances in networking, wireless communication, mobile devices, and sensor networks, have resulted in a proliferation of novel protocols for accessing and distributing information. These protocols often differ in how users specify the intended recipients of a message. Some protocols, such as IP unicast, use a unique name for each destination node, while others, such as geocast, use geographic location to specify message recipients. Recent research in sensor networks and mobile computing have resulted in specialized protocols that allow an application to specify the destination using a combination of constraints over name, space, and time [3] [1] [7]. Despite the number and variety of communication protocols, no framework is available for unifying the name, time, and space domains. Such a framework would be useful for formally characterizing existing protocols, evaluating their relationships to each other, and serving as a basis for discovering new protocols. This section presents a notation that creates such a framework.

In order to create a unified destination specification framework, we need to abstract away the specifics of a particular protocol and focus on higher-level con-

cepts. Thus, instead of developing a notation based on protocol-specific attributes like IP addresses or GPS coordinates, we simply use abstract notions of identity (i), physical location (l), and time (t). To formally define the domains of these variables, we use I , S , and T to be the universal spaces of identity, space, and time, respectively. We assume that a node must be in a particular location, $l \in S$, at a particular time, $t \in T$. We also assume that all nodes are distinguishable and thus, at a high level, have unique IDs. Whether they are actually assigned unique IDs, however, is dependent on the needs of the protocols. For those protocols that require unique IDs, such as IP unicast, each node must be assigned an ID. For those protocols that do not require unique IDs, such as SPEED or directed diffusion, the nodes need not be assigned IDs.

So far, we have described three attributes of a node that can be used for specifying the destination of a message. Next, we turn our attention to the dynamics of a node with respect to these attributes. Since a node represents a computational device that may be turned on and off, the node may join and leave the network over time. We assume that a node has a unique ID that remains static regardless of its status. We treat a node as non-existent when it is off. Over time, the node may change its location, especially if it is a mobile device such as a laptop or PDA.

Imagine a three-dimensional graph whose axes represent identity, space, and time. At any time, each node is represented by a point in the graph. Over time, a node is represented by a sequence of dots that collectively reflect its movement. This sequence of points forms the *movement profile* of the node. A delivery specification can then be visualized as a subset of points called the *delivery volume* such that if a nodes' movement profile intersects the delivery volume, it is delivered the message. More precisely, if a node with identity i , henceforth called "node i ," enters the delivery volume, it will be delivered the message while it is still in the delivery volume. For this purpose, we introduce three predicates, ω , ψ , and ϕ , that operate over identity, space, and time, respectively. The following notation specifies the nodes that should receive message m :

$$m : \langle i, l, t :: \omega(i), \psi(l), \phi(t) \rangle \quad (1)$$

The three predicates collectively determine the set of destination nodes. A node is a destination if its identity satisfies ω , its location satisfies ψ , and the current time satisfies ϕ . We use predicates to capture the delivery specification, as opposed to explicit sets, to allow for generalizations across dimensions. For example, one generalization is $\omega(i, t)$ in which the targeted identities

change over time. These generalizations are explored further in the next section.

The notation above is interpreted as follows: every destination node will receive the message at least once while it is still in the delivery volume and no other nodes will receive the message. The specification is strong in that all destination nodes receive the message and no other nodes receive it.

We formalize our notation by treating it as a relation between the action history of the nodes and delivery history of the message. Formally, a history set, H_i , is associated with node i . It records all actions performed by i . An entry in H_i is denoted as (i, l, t, α) , where α is an action such as $\text{received}(m)$. Thus, if $(i, l, t, \text{received}(m)) \in H_i$, node i received message m while located at l and at time t . For convenience, we assume that for each $t \in T$, there exists an entry in H_i . If the node did not perform any action at that time, then a *null* action is used as a place holder. The time component of the history provides a total ordering of all the actions taken by the individual nodes.

We use the notation $[i, l, t]$ to indicate the occurrence of an event on node i located at l at time t . That is,

$$[i, l, t] \equiv \langle \exists \alpha :: (i, l, t, \alpha) \in H_i \rangle \quad (2)$$

This serves as an accessor to the action history of a node.

We use the notation $m : [[i, l, t]]$ to capture a delivery of message m . That is,

$$m : [[i, l, t]] \equiv (i, l, t, \text{received}(m)) \in H_i \quad (3)$$

It states that message m was received by a node with identity i , at location l , and at time t .

Using the above notations for a node's action history and message delivery, we formally define a message delivery specification, $m : \langle i, l, t :: \omega(i), \psi(l), \phi(t) \rangle$, as follows:

$$\begin{aligned} m : \langle i, l, t :: \omega(i), \psi(l), \phi(t) \rangle &\equiv \\ &\langle \forall i : i \in I :: \\ &\quad \langle \exists \lambda, \tau : \omega(i) \wedge \psi(\lambda) \wedge \phi(\tau) \wedge [i, \lambda, \tau] :: \\ &\quad \quad \langle \exists \lambda', \tau' : m : [[i, \lambda', \tau']] :: \\ &\quad \quad \quad \psi(\lambda') \wedge \phi(\tau') \rangle \rangle \rangle^1 \end{aligned}$$

¹The three-part notation $\langle \mathbf{op} \text{ quantified_variable} : \text{range} :: \text{expression} \rangle$ used throughout the text is defined as follows: The variables from *quantified_variables* take on all possible values permitted by *range*. If *range* is missing, the first colon is omitted and the domain of the variables is restricted by context. Each such instantiation of the variables is substituted in *expression*, producing a multiset of values to which **op** is applied, yielding

This formalization reads as follows: “For all identities, i , if node i was in the delivery volume, it was delivered the message while still in the delivery volume.”

There are two alternative specifications that we considered. The first assumes best effort delivery. That is, there is no guarantee that every destination node will receive the message, or that delivery will occur only once. It is also possible for nodes outside the delivery volume to receive the message. This interpretation is nice in that it does not force protocols to provide strong guarantees, and thus is more likely to be implementable. In fact, most multicast protocols, such as IP multicast, directed diffusion, geocast and mobicast [3], fit this interpretation exactly. However, its lack of guarantees and preciseness make it less useful as a specification language.

Another, stronger, alternative assumes that all nodes in the delivery volume receive the message exactly once, and no other nodes receive the message. This is attractive because of its preciseness, which simplifies the analysis of higher-level protocols. The problem, however, lies with the difficulties involved in achieving such semantics. In general, a lot of machinery must be put in place to ensure a single delivery to all destination nodes, which results in a lot of computational and bandwidth overhead. For example, a tree structure that is rooted at the sender and touches all receiver nodes may be used. But creating and maintaining such a tree is difficult, particularly in dynamic environments where the topology of the network constantly changes. For this reason, many existing protocols such as IP multicast, directed diffusion, and geocast, only offer best-effort delivery. It is possible for the same message to be delivered to a particular node multiple times due to topology changes or message flooding. Thus, for lack of practicality, this interpretation is also undesirable.

We have, at this point, adopted a specification language that balances preciseness with feasibility, but is slightly idealized in that it assumes an ideal protocol and a failure-free network. In a real network, a message may be lost, delayed, or delivered too early. Sometimes, a message may be duplicated during delivery, resulting in a node receiving the same message twice. While some of them, such as early delivery and duplicate detection, can be easily addressed using queues and history, they needlessly complicate the delivery semantics. We decided to adopt an idealized approach because it elegantly captures the ultimate goal of the

the value of the three-part expression. If no instantiation of the variables satisfies *range*, the value of the three-part expression is the identity element for **op**, e.g., *true* when **op** is \forall or zero if **op** is “+”.

delivery protocols. The next section provides generalizations of our specification. A discussion on extensions to our notation is given in section 3.5.

We now discuss how our notation captures IP unicast. Other protocols, like geocast, ComMotion, and Mobicast, will be discussed in the next section. In IP unicast, the sender explicitly identifies the recipient using the recipient’s unique IP address. Assuming the destination node’s IP address is ip_0 , this is captured as follows:

$$m : \langle i, l, t :: i = ip_0, l \in S, t \in T \rangle \quad (4)$$

Since $l \in S$ and $t \in T$ both evaluate to *true*, this simplifies to:

$$m : \langle i, l, t :: i = ip_0, true, true \rangle \quad (5)$$

This specification represents IP unicast since it restricts the recipient based on identity and does not impose any restrictions on the spatial or temporal dimensions. Note that since there is no restriction on the temporal dimension, all unicast operations last forever. That is, if a message is sent at time t_s , any node with a name i that turns on after t_s will receive the message regardless of how much time has passed. This does not match the real behavior of IP unicast. We address this anomaly by introducing the timing constraint $\phi(t) = (t \in [now, now + \delta])$, which specifies a bound on the time to deliver the message. By and large, the choice of δ can sure that the message is delivered to all destination nodes that existed at the time the message was sent.

In this section, we introduced our notation for specifying the destination nodes of a message, provided a formal definition of the notation, and gave an example of how it is used to express IP unicast. In the next section, we provide a more in-depth investigation of how our notation can be used and generalized to express existing communication protocols and reveal interesting future directions.

3 Transport Protocols Revisited

This section shows how our specification notation captures a variety of popular communication protocols. We categorize them based on the primary dimension they use for specifying the destination nodes. These include identity-based protocols, location-based protocols, and time-based protocols. We also present protocols that take multiple dimensions into account and show how our notation can be generalized to account for interdependencies among dimensions. The section

ends with a discussion on further generalization that can enhance our notation.

3.1 Identity-Based Protocols

Traditionally, transport protocols have been built around the concept of identity. This subsection analyzes broadcast and multicast as well as certain application-specific protocols. Of great interest is how to characterize the address of a host. An address may refer to a single host, as is the case for unicast, or a group of hosts, as is the case for multicast and broadcast. A multicast address can be modelled as a set defined by enumerating the identities of each node that must receive the broadcasted message. However, if every host is characterized by certain attributes, we can define the destination nodes in terms of their attributes. We generalize our notation to allow enumerating the destination nodes based on their attributes rather than identities and an example application is provided.

Broadcast. Broadcast is a one-to-all communication paradigm. We capture broadcast’s specification as follows:

$$m : \langle i, l, t :: true, true, true \rangle \quad (6)$$

The specification of broadcast does not impose any constraints on the delivery of messages.

Broadcast’s specification, while conceptually simple, raises the important issue of what is the precise semantic meaning of broadcast. The problem is exacerbated in a dynamic environment like an ad-hoc network. If a node is part of the network for only a few seconds should it receive the message? We feel that defining strict semantics using formal methods is necessary to disambiguate the meaning of communication protocols. Accordingly, we have introduced in the previous section, formal semantics for our notation. Our notation specifies that all nodes in the delivery volume must receive a message before leaving the delivery volume.

In order to be able to provide efficient implementation to the broadcast protocol, we need to acknowledge the fact that it may be difficult, if not impossible, to implement the idealized definition presented above. To alleviate this problem, we can associate with the broadcasted messages a deadline. In this case, the new definition is:

$$m : \langle i, l, t :: true, true, t \in [now, now + \delta] \rangle \quad (7)$$

Multicast. Multicast is a one-to-many communication paradigm involving a source node that sends

messages to a group of nodes. Multicast is commonly used by applications such as video conferencing and internet radio. A prerequisite for communication is to identify the group of destination nodes. Formally, this requirement is captured by defining a predicate over the identity domain as follows:

$$m : \langle i, l, t :: \omega_\alpha(i), true, true \rangle \quad (8)$$

where ω_α is the characteristic function of the multicast address α i.e. ω specifies what nodes should receive the multicast message. The above formalism insulates us from explicitly managing the membership of nodes in the multicast group. In IP multicast, nodes are added to the group using the join operation and nodes, are removed from the group using the leave operation. As the group membership changes over time messages will be delivered to the current members of the group as characterized by ω_α . Note that the application does not need to concern itself with these changes.

Profiles. An alternative to fixed naming is attribute-based naming. It is commonly used in ubiquitous computing, peer-to-peer computing, and sensor attributes. In this context, each node is described in terms of its properties. The aggregation of attributes forms a profile that describes the node and may vary in time. For instance, a temperature sensor can be characterized by its available power. This may allow the specification of a protocol that can adapt to availability of energy to minimize energy consumption

An application may take interest in a particular group of nodes selected based on their properties. To achieve this, the attribute-based naming is pushed to the protocol level, allowing direct communication with a set of nodes created based on their attributes. Such an approach is introduced in [2].

Formally, we can capture the process of selecting a group of nodes based on their attributes by defining a predicate that holds true for the desired set of nodes. To simplify the notation, we define P to be the set of all profiles and redefine the identity domain to be $I \equiv \zeta \times P$, where ζ is the set of identifiers. Previously, we used $I = \zeta$ and we will return to this simpler interpretation in the remainder of the section.

Let’s consider a sensor network that gathers data for a weather predicting application. For this application we use three types of sensors: temperature, pressure, and wind direction. The notation allows us to select sensors of a particular type.

By introducing the concept of sensor type, we can preferentially communicate with a specific type of sensor. Notice the subtle transition from the idea of explicitly enumerating the members of a group based on their address to implicitly defining (generating) a group

based on their properties. Already, research in how to meaningfully take advantage of this mechanism is underway, initial results being presented in [4]. Formally, we can capture a communication protocol for these applications as follows:

$$m : \langle i, l, t :: \omega_{type}(i), true, true \rangle \quad (9)$$

where

$$\omega_{type}(i) \equiv (i.type = temperature)$$

assuming the *type* property within *i*'s profile is accessed by *i.type*.

To summarize, we formally captured broadcast and multicast, and analyzed the impact of fixed-naming and attribute-based naming on identifying the intended message destinations.

3.2 Location-Based Protocols

Traditionally, nodes are assumed not to be aware of this physical location. However, in sensor networks, location awareness is a common assumption in protocol design. A sensor network is a self-organizing network potentially composed of thousands of miniature computational devices with sensors that are scattered throughout the environment. Many new protocols rely on location-awareness [4] [2] [8] [3] [6] [5].

Geocast. Geocast [9] delivers a message to all nodes located in a certain geographic area. It is defined as follows:

$$m : \langle i, l, t :: true, \psi(l), true \rangle \quad (10)$$

where $\psi(l)$ characterizes the delivery area. Geocast has many applications. For example, a weather advisory system may deliver warnings to a particular area. Suppose this area is a circle with a radius r that is centered at coordinates (c_x, c_y) . Assuming a nodes' location, l , is characterized by an xy coordinate accessed by $l.x$ and $l.y$, $\psi(l)$ is defined as:

$$\psi_{c_x, c_y, r}(l) \equiv ((c_x - l.x)^2 + (c_y - l.y)^2 \leq r^2)$$

where the subscripts indicate application-defined parameters. Note that ψ evaluates to *true* if the node is located in the circle.

CommMotion. CommMotion [8] is a system that allows users to receive messages only when they are located at particular points in space. For instance, a user may create a grocery list and expect to receive it when he enters the grocery shop. Assuming the geographic location of the grocery shop is characterized by

the same circle used in Geocast, CommMotion can be specified as follows:

$$m : \langle i, l, t :: \omega(i), \psi(l), true \rangle \quad (11)$$

where $\omega(i)$ specifies which nodes need to receive the message and $\psi(l)$ characterizes the location where the nodes must be to receive it.

In the above example, delivering the shopping list when the user is in the grocery can be modelled as follows. The user is connected using a node with identity $i = ip_0$. The area of the grocery shop is characterized by a circle of radius r and centered at the origin. We can define ω :

$$\omega(i) \equiv (i = ip_0)$$

where ω is *true* when the evaluated node is identified by ip_0 . ψ is defined by the same formula as in the case for Geocast and is *true* when node ip_0 is located in the circle.

Though the specifications of Geocast and CommMotion may seem similar, they differ in the way message delivery occurs. Geocast specifies the fact that we want to deliver to a particular delivery zone and any node that may be present in the delivery zone will receive the message. In contrast, CommMotion says that the message is delivered to a zone and a user receives it when it enters the zone.

Geocast is a representative for location-aware protocols. CommMotion has some interesting features as it couples the location of the user with message delivery. The two examples are intended to provide a flavor of how to characterize the spatial dimension and how it may be correlated to other dimensions, as is the case with CommMotion.

3.3 Time-based protocols

Real-time systems are concerned with making the messages delivery time predictable and bounded. Two types of semantics can be identified in characterizing the time of delivery: deadline based and just-in-time. To illustrate both cases let's reconsider the example of the weather advisory system that was defined as:

$$m : \langle i, l, t :: true, \psi(l), true \rangle \quad (12)$$

To transform this specification into one that involves a deadline we add a predicate on the temporal dimension:

$$m : \langle i, l, t :: true, \psi(l), t \in [now, now + \delta] \rangle \quad (13)$$

where δ is the deadline and *now* refers to the time when the message was sent. We can change the semantics of the protocol to just-in-time delivery as follows:

$$m : \langle i, l, t :: \text{true}, \psi(l), t \in [t_1, t_2] \rangle \quad (14)$$

where $t_1 = \text{now} - t_0 + \varepsilon$ and $t_2 = \text{now} + t_0 + \varepsilon$. The delivery has to occur in a tight interval 2ε which formally captures the concept of just-in-time delivery.

SPEED. SPEED[1] is a novel soft real-time communication protocol for sensor networks that focuses on making the delivery of messages predictable. To this end, SPEED guarantees uniform delivery speed across the network. Since we are considering a sensor network, it is reasonable to assume that nodes are location aware. As such, the location of the node can be defined by its cartesian coordinates. When the speed of message delivery and the distance between nodes are known, the deadline for end-to-end message delivery is:

$$t_{\text{end-to-end}} = \frac{\text{dist}(i_0, i_1)}{v} \quad (15)$$

The distance is the Euclidian distance between the source of the message, node i_0 , and the destination of the message, node i_1 ; v is the velocity of message propagation through the network. The delivery of messages to a destination can be formalized as follows:

$$m : \langle i, l, t :: i = i_1, \text{true}, \phi(t) \rangle \quad (16)$$

where ϕ imposes the SPEED's realtime constraints on message delivery and is defined as:

$$\phi_{i, i_0, v, \varepsilon}(t) \equiv (t \in \frac{\text{dist}(i, i_0)}{v} \pm \varepsilon)$$

Thus, ϕ is *true* if the end-to-end delivery time is within a tight interval having the expected end-to-end delay given by equation 15. ε is the tolerance of the system relative to the tardiness in the message delivery.

Time-based anycast. Often, systems do not have control over the delays incurred during message transmission. In this case, the applications must adapt its delivery behavior to account for this uncertainty. The following protocol can be seen as a transport protocol adaptation to delays incurred during transmission. Consider the scenario of having to deliver a letter to an office. The letter can be received by any of the available secretaries, A, B, or C. Assuming that each secretary has its own schedule, modelled using a function $\xi : I \rightarrow R$ that associates secretaries and schedules. Two secretaries are never at work simultaneously. A protocol specification for the above example can be formalized as follows:

$$m : \langle i, l, t :: \omega(i, t), \text{true}, \text{true} \rangle \quad (17)$$

where $\omega(i, t)$ is used to correlate the identity and the schedules as follows:

$$\omega(i, t) \equiv (i \in \xi^{-1}(t))$$

where ξ^{-1} , the inverse of ξ , identifies which secretaries are available at a given time. The delivery can be accepted by any available secretary.

3.4 Spatiotemporal protocols

With the advent of sensor networks and mobile computing, a new generation of spatiotemporal protocols have been developed. One example is Mobicast [3].

Mobicast. Mobicast is a multicast protocol that allows a node to send a message to all nodes in a geographic region that changes with time. This region is called the *delivery zone*. Ideally, all nodes traversed by the delivery zone will have received the message.

Mobicast is useful in many scenarios. Consider an ambulance speeding down the road. To avoid slowing down or getting into an accident, it needs to warn cars ahead that it is about to come. This warning message can be broadcasted to cars ahead of the ambulance. However, the message must not be delivered too far ahead of the ambulance. Otherwise, cars may pull off the road and wait a long time before the ambulance actually passes. Mobicast may be used to control the propagation of the warning message ahead of the ambulance. The delivery zone is specified such that it moves ahead of the ambulance an appropriate distance.

We formally characterize Mobicast as follows. The delivery zone is modelled using $Z(t)$. $Z(t)$ maps the time to the geographic region a message must be delivered to. Once a node enters $Z(t)$, it must be delivered the message as soon as possible. Let t_s be the time in which the message is sent and t_d be the lifetime of the message. The delivery specification of Mobicast is:

$$m : \langle i, l, t :: \text{true}, \psi(l, t), t \in [t_s, t_d] \rangle \quad (18)$$

where

$$\psi(l, t) \equiv (l \in Z(t))$$

Notice that ψ correlates location with time. This is necessary to capture a delivery zone that changes over time. Mobicast is only the first among what we expect to be a large family of new spatiotemporal protocols.

3.5 Further Generalizations

Our current notation requires a message be delivered to *all* nodes within the delivery volume. This is

not always feasible or desirable. For example, if, while a message is being propagated, a node quickly moves in and out of the delivery zone, the protocol implementation might not have a chance to deliver the message to it. A simple way to account for this is to refine the specification to include a minimum time, T_{min} , that a node must be in the delivery volume in order for delivery to be required.

Another area where our notation may be refined is the ability to specify how many destination nodes actually receive the message. Currently, we have defined the specification to mean that *all* nodes in the delivery volume will receive the message. Some applications may require only one node within the delivery volume to receive the message, so called *anycast*. This has immediate applicability to many scenarios. For example, suppose we have a sensor network that measures the ambient temperature of a field. An observer is interested in the temperature of a particular region of the field. This region is small enough such that temperature fluctuations across it are minimal. Thus, when the observer queries the region for its temperature, only one node in the region needs to reply to the query. Having all nodes in the region receive the query increases network congestion and also wastes valuable battery power.

Anycast is only one alternative mode of delivery, there are many other possibilities. For example, some applications may only require that *most* nodes in the delivery volume receive the message or that only nodes whose battery charge level is at least 90% receive the message. The possibilities are endless.

4. Conclusions

Computer networks revolutionized the way we communicate and live our lives. Recent advances in sensor networks, mobile computing, and wireless communication, promise to continue to shape the way we interact with each other and with our environment. The advent of these new computing paradigms resulted in a plethora of new information dissemination protocols. While traditional protocols assigned each node a unique identifier that is used as the destination address of a message, new protocols tend to rely on an interplay among multiple dimensions including space and time. We observed that the various protocols differ mainly in the way they specify which nodes a message must receive a message. These differences were, at times, subtle, and relationships between protocols unclear. Given this observation, we developed a new notation for unifying the three specification dimensions of identity, space, and time. We formally defined the

notation and showed how it captures the delivery semantics of popular protocols in use today. We also showed how our notation can be generalized to capture new protocols that have yet to be investigated and discussed further enhancements to our notation that allow it to capture entirely new classes of protocols leading to new avenues of research.

Acknowledgements. This research was supported in part by the National Science Foundation under grant No. CCR-0325529 and by the Office of Naval Research under 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 research sponsors.

References

- [1] T. He, J. Stankovic, C. Lu, and T. Abdelzaher. Speed: A stateless protocol for real-time communication in sensor networks. In *International Conference on Distributed Computing Systems (ICDCS 2003)*, 2003.
- [2] J. S. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, and D. Ganesan. Building efficient wireless sensor networks with low-level naming. In *Symposium on Operating Systems Principles*, pages 146–159, 2001.
- [3] Q. Huang, C. Lu, and G.-C. Roman. Spatiotemporal multicast in sensor networks. In *First ACM Conference on Embedded Networked Sensor Systems (SenSys'03)*.
- [4] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *Mobile Computing and Networking*, pages 56–67, 2000.
- [5] X. Jiang and T. Camp. Review of geocasting protocols for a mobile ad hoc network. In *Proceedings of the Grace Hopper Celebration (GHC)*, 2002.
- [6] Y. Ko and N. Vaidya. Geocasting in mobile ad hoc networks: Location-based multicast algorithms. In *Second IEEE Workshop on Mobile Computer Systems and Applications*, 1998.
- [7] C. Lu, B. M. Blum, T. F. Abdelzaher, J. A. Stankovic, and T. He. Rap: A real-time communication architecture for large-scale wireless sensor networks. In *Eighth IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'02)*, 2002.
- [8] N. Marmasse and C. Schmandt. Location-aware information delivery with commotion. In *HUC*, pages 157–171, 2000.
- [9] J. C. Navas and T. Imielinski. Geocast - geographic addressing and routing. In *Mobile Computing and Networking*, pages 66–76, 1997.