

UC San Diego

Technical Reports

Title

Fulcrum -- An Open-Implementation Approach to Context-Aware Publish / Subscribe

Permalink

<https://escholarship.org/uc/item/7sb3s217>

Authors

Boyer, Robert T
Griswold, William G

Publication Date

2004-06-08

Peer reviewed

Fulcrum – An Open-Implementation Approach to Context-Aware Publish / Subscribe

Robert T. Boyer William Griswold
University of California, San Diego
La Jolla, CA 92039-0114
+011.858.442.1322
{ rboyer, wgg }@cs.ucsd.edu

ABSTRACT

Content-based publish / subscribe (CBPS) systems are a natural substrate for context-aware applications because they provide the right separation of concerns, efficient event distribution, extensibility, and scalability. However, the separation of concerns afforded by CBPS middleware precludes publishers from collaborating with each other to efficiently detect context-aware conditions for publication as events. We overcome this problem with an open implementation approach, which enables subscribers to attach a domain-specific *implementation strategy* to their context-aware subscriptions. Strategies are supported by first-class *active subscriptions* that can deploy lower-order dynamic publisher / publisher cross-subscriptions where events enter the CBPS system. By exchanging data on an as-needed basis, event traffic can be dramatically reduced. Our algorithm for detecting the proximity of mobile buddies reduces event traffic from $O(|\text{events}|)$ to $O(|\text{movement}|)$ worst case, $O(\lg|\text{movement}|)$ expected. Experiments reveal an 82% reduction with current positioning technologies.

Categories and Subject Descriptors

D.2.11 [Software Engineering]: Software Architectures – *distributed processing, separation of concerns*. D.3.3 [Distributed Systems]: Publish-subscribe.

General Terms

Performance, Design, Languages.

Keywords

Context-aware computing, content-based publish / subscribe, distributed event-based systems.

1. INTRODUCTION

Context-aware computing takes data from a highly dynamic environment and synthesizes it into information and knowledge for decision-making. Of particular interest are the relationships between changing data. Simple examples include: your location with respect to mine, when my objective is to speak with you in an impromptu hallway meeting; or stock price with respect to earnings, i.e., P/E ratio, when the task is to increase retirement investments. Specific context of interest depends on individual users and their current tasks and objectives. Tarasewich provides a useful context-awareness survey [14].

In a sensor network composed of low-powered components deployed in an ecological preserve, we might tag animals of interest with a battery-powered GPS receiver, CPU, and wireless network-

ing. Scientists wish to monitor, throughout the year, the interactions between various animals. To preserve power, communications need to be minimized. The scientists are initially interested to learn when the proximity of certain animals is less than 10 meters (e.g., $(X_1 - X_2)^2 + (Y_1 - Y_2)^2 < 10^2$). As these relationships are observed, the system must be able to support measurement for new hypotheses that are generated by the scientists. Also, multiple researchers may observe the same environment with different goals. For reasons of cost and feasibility, the sensor network's resources need to be shared.

Generally, we can anticipate that contextual data will ultimately derive from “billions of users connected to millions of services using trillions of devices” [5]. Myriad sensors reporting potentially interesting raw data at high data rates will be as widely distributed as an equally diverse set of consumers, who may themselves be information providers. To effectively support context-aware computing in this environment requires efficient, extensible, and scalable data distribution and processing.

Content-based publish / subscribe (CBPS) systems hold the potential for such a solution. In publish / subscribe systems, publishers publish *events* and subscribers subscribe to *events*, with the CBPS middleware doing event matching and routing. Hence, publishers and subscribers do not need to know about each other, enabling new publishers, subscribers, event types, and subscriptions to be freely added to the milieu.

Not only does the middleware provide separation, but also efficiency: it filters out new events against subscriptions at the publisher's event broker, exploits overlapping subscriptions, and employs multicast-like routing of events to subscribers. Efficient filtering at the publisher's event broker is achieved by content-based pattern matching against a publisher's event in a series of independent filters, e.g., $\{(e.x < 10) \& (e.y > 30)\}$. Sequences of events can be similarly pattern matched [1,2]. Context-aware relationships like the proximity relationship entail comparing event attributes to each other. This requires subscribers to subscribe to raw location and compute the distance themselves, and the efficiencies of evaluating at the publisher's event broker are lost.

Recent work enables the aggregation of attributes from multiple data streams with more complex processing and filtering performed within the network [3,11]. Common aggregations and transformations can also be shared [3]. It is possible, then, to evaluate the proximity relationship at the first common event broker node ($1^{st}CN$) that connects the publishers with the subscriber.

When publishers are distant from the $1^{st}CN$, each intervening event broker node must process and forward all events, which is costly. To get the best possible performance requires evaluating

context-aware relationships at the publisher's event broker. To do so efficiently requires knowledge of the modeled relationship as well as how it will be used. Only the subscriber has this domain knowledge, and has no way of sharing it with the middleware.

The open implementation software design technique [12,13] was developed for just such situations. In the open implementation approach, a module's interface is designed to allow a client to assist in the selection of the module's implementation strategy. The module's auxiliary interface may allow a client to describe its usage patterns, to specify an implementation (e.g., hash table), or even to provide its own implementation – adhering to well defined interface specifications. This allows the client to tailor the module's implementation strategy to better suit their needs, while retaining the advantages of closed implementation modules (i.e., the traditional black box).

Following the open implementation paradigm, Fulcrum creates an efficient substrate for context-aware publish / subscribe (CAPS). For the example of tracking deer proximity, Fulcrum enables a reduction in the event traffic from $O(|\text{events}|)$ to $O(|g_2|\text{movement}|)$ expected, using domain-specific optimizations that filter events at the publishers' event brokers (called *entry nodes*). Fulcrum provides this facility as a three-part extension to CBPS while preserving its anonymous, asynchronous, and loosely coupled nature:

1. A subscriber can subscribe directly to a *derived relationship*, such as the distance between two entities. A derived relationship has no publisher, *per se*.
2. The derived relationship is realized through *active subscription* support. An active subscription, like a Solar operator [3], is a first-class pub/sub Java applet; it can be both a publisher and subscriber. In particular, an active subscription is responsible for:
 - a. subscribing to the event data comprising the relationship;
 - b. instantiating a publisher of the derived relationship.
3. A separate *implementation strategy* can be attached to the relationship subscription to exploit the domain's semantics to improve efficiency. In particular, it can push the subscription to the entry nodes and establish a secondary set of subscriptions and publications that permit the relationship to be computed in a distributed, collaborative fashion, thus reducing event traffic.

Note that the same property can be implemented by different strategies as appropriate to the context of use. Similarly, the same implementation strategy can be reused among properties with similar semantics (e.g., the notion of distance).

The remainder of this paper is organized as follows. Section 2 discusses the state of the art capabilities of CBPS and their applicability to context-aware computing. Section 3 describes our design of Fulcrum. Using buddy proximity as an example, Section 4 details how an implementation strategy is deployed to achieve efficiencies. Section 5 evaluates the system as built. Section 6 discusses the tradeoffs with our approach, and Section 7 concludes.

2. CONTENT-BASED PUB / SUB

In the current state of the practice, as exemplified by Siena [1,2], CBPS systems are composed of three components. One, a *publisher* provides *events (messages)*, each in the form of a tuple sequence $\{(name_1, type_1, value_1), (name_2, type_2, value_2), \dots\}$. Two, a *subscriber* requests events of interest by using *subscription filters* of the form $\{(name_1, operator_1, value_1), (name_2, operator_2, value_2), \dots\}$, where each *operator* will be a relational operation like $<$, $=$, etc. Three, *event brokers* mediate between publishers and subscribers, providing an application-level overlay network for efficiently matching and routing events, providing independence of publishers and subscribers.

In particular, publishers *advertise* their event *types* with brokers, and subscribers register their interest in events through their content subscription filters. Abstractly, the brokers then check every published event against every filter, passing on those events that satisfy a filter. Concretely, the brokers set up a "switching fabric" between publishers and subscribers by pre-matching filters to the advertised event types (i.e., a publisher's event type promises to contain all the required names of a subscriber's filter).

This switching fabric has two performance benefits for the brokers. One, filters are only applied to events that have a chance of satisfying the filter. Two, it is possible to avoid redundant filtering for overlapping subscriptions, as well as send only one copy of each event between brokers when multiple subscribers share parts of the same pathway from the publisher to an intermediate broker. Subscription filters are pushed upstream from the subscriber toward the information producer to quickly suppress information at broker nodes for which no downstream subscriber is currently interested. This eliminates unnecessary network traffic and excess computation at internal event brokers and at the end client. Consequently, the client application is, in effect, being pushed into the network via the middleware of the CBPS system.

CBPS systems seem ideal for supporting extensible, scalable context-aware systems. The brokered publish/subscribe paradigm makes publishers and subscribers largely independent from each other, while providing economies of scale through sharing in the overlay network. Publishers and subscribers can be readily added, accommodating new application functionalities, and new event brokers can be added to improve the efficiency of the network, thus providing Internet-level scalability [1,2].

2.1 Context-Awareness Requirements

At its core, a context-aware system needs to detect and react to *situations*, that is, the moment-to-moment circumstances of the entities it supports, tracks, or models. This includes awareness of the changing *relationships* amongst entities. In exploring the decentralization of ActiveCampus into a federation of cooperating peer servers, our first challenge was efficiently detecting relationships amongst users. In the federated model, each ActiveCampus user would be logged into a "home" server that is responsible for reporting (publishing) that user's context, thus making CBPS a natural substrate. An example relationship would be when a person moves into the vicinity of a buddy, which should result in notifying them of their proximity. This might motivate either person to seek the other out or contact them by instant messaging [8]. A more general example would be the ability of a user to detect when a critical mass of friends are gathering at a location.

The buddy relationship situation is naturally described with a formula¹ like $(X_R - X_W)^2 + (Y_R - Y_W)^2 < D^2$. However, the core capabilities of CBPS do not allow for a subscription to compare event attribute values coming from two different publishers. As a consequence, the subscriber must issue separate subscriptions for Robert's location and William's location. This results in their location events being pushed all the way through the network to the subscriber, which computes their distance itself, likely learning that the friends are no where near each other. Network cycles and bandwidth and subscriber cycles are wasted.

To directly and efficiently support the detection and reporting of relationships, then, a CBPS system must:

1. Enable intra- and inter-event attribute-to-attribute computations (e.g., to compare X and Y attributes from Robert's position event with X and Y attributes from William's position event).
2. Enable a new event type to be advertised and subsequently published from within the CBPS middleware. Such an event either aggregates attributes from other "sub-events" into a new unified event or abstractly represents the satisfaction of a relationship subscription like $(X_R - X_W)^2 + (Y_R - Y_W)^2 < D^2$ in a form like "Robert is close to William."
3. Maintain state information at event brokers to support data aggregation. In particular, two events contributing to an aggregate event (e.g., William's new position and Robert's new position) may not arrive at the same time, so an event broker will need to store events until all the required events are present.
4. Suppress the propagation of sub-events at the publisher's event broker if it could not satisfy an aggregate event. Suppressing at a common node can be wasteful for many broker network configurations. For example, if Robert and William are far apart, and moving slowly, it would be best if their respective position reports were not pushed through the broker network.
5. Provide mechanisms to atomically query on, and subscribe for, an event. Akin to fetch-and-add, such a feature prevents race conditions when the contents of a subscription depend on dynamic data. For example, in adjusting a "suppression limit" based on an event's current attribute value, the value could change between the time of query and time of subscription, causing the desired event to not be propagated, thus missing a key state transition.

Of course, in satisfying these requirements we desire that the transparency and separation of concerns provided by CBPS, such as the decoupling of producers and consumers, be retained.

2.2 CBPS State of the Art – Related Work

A growing number of CBPS systems are being developed with extensions on the core set of capabilities described above. Gryphon [11] and Solar [3] are exemplary.

Gryphon aims to improve overall system performance of CBPS with *stateful, relational subscriptions* where the middleware becomes responsible for computing derived state instead of the end client [11]. Relational subscriptions enable aggregating events from multiple publishers in the form of a composite event. Such an event can be filtered based on the attributes of two publishers

¹ Subscripts are used here to denote separate events from possibly different publishers.

$\{(X_1 < 10) \text{ and } (Y_2 > 30)\}$, thereby suppressing the event at the 1stCN.

Their approach is akin to a database view across multiple tables, say from multiple publishers. A "publisher topic is a 'source' relation representing an event history, where each tuple corresponds to an event." A "subscription is a request to receive incremental updates to a derived view defined by a relational expression on one or more base relations or other views." The capabilities for the SQL-like behavior are directly built into the overlay network. The middleware capabilities "incrementally maintain the states and deliver the updates." [11]

In the Gryphon approach, the proximity of Robert and William would be computed by specifying a join of their location views at the 1stCN, followed by a select on the one-row table with a predicate like $\{(X_R - X_W)^2 + (Y_R - Y_W)^2 < D^2\}$. The authors acknowledge that squelching at the 1stCN is inadequate, and proffer an idea they are working on called "selective curiosity", which makes the aggregation node responsible for pushing event reduction clues back to the information providers.

Solar extends CBPS with operators to filter, aggregate, or transform event data [3]. "Applications describe their desired event stream as a tree of operators that aggregate low-level context information published by existing sources into the high-level context information needed by the application." [4] The Solar middleware, with the help of a centralized manager node, deploys a subscription in the form of Java operator objects, placing the objects to maximize sharing subscribers and minimize network traffic (i.e., placement at 1stCNs). In essence, the result is an acyclic dataflow program. In the Solar approach, the proximity of Robert and William would be determined with operators that aggregate their location events into a single event at the 1stCN, transform it into distance, and then apply a filter for the distance constraint.

Gryphon and Solar provide essential characteristics for context-aware computing, both increasing expressiveness and increasing efficiency through operator re-use. With either of these systems, subscribing to the proximity relationship requires the 1stCN to separately subscribe to all position reports from each subject of interest, $\{\text{NAME}=\text{Robert and } X=\text{ANY and } Y=\text{ANY}\}$ and $\{\text{NAME}=\text{William and } X=\text{ANY and } Y=\text{ANY}\}$ in order to aggregate, evaluate, and report on the satisfaction of the relationship. As shown in Figure 1, the location reports are passed through and processed at every intermediate event broker until the common




Figure 1. Relationship detection requires all data to pass through all event brokers until a common node is reached. Letters indicate processing occurring at the event brokers. Diamonds are used to represent publishers; circles are event brokers; and hexagons are subscribers or subscriber based processing. The solid lines represent overlay network connectivity while the dashed arrows represent event movement.

node receives and processes the event to determine satisfaction of the relationship. This succeeds in off-loading the subscriber, but still loads the network with $O(|\text{events}|)$ event traffic and processing at every intervening node between the publishers and the 1stCN, and additionally burdens the 1stCN with frequent (and perhaps fruitless) proximity computations.

3. FULCRUM

“Give me a place to stand and a lever long enough and a fulcrum on which to place it, and I shall move the world.” – Archimedes, 220 B.C.

These costs are unacceptable in context-aware environments, where there can be thousands of publishers (devices) publishing events at high rates. Ideally, the middleware would permit suppressing location events at their entry nodes, only forwarding those that could satisfy a derived relationship event.

Fulcrum is a context-aware publish / subscribe (CAPS) system that employs open implementation to give subscribers the ability to efficiently and transparently subscribe to relationships. That is, in addition to the ability to subscribe to relationships like Gryphon and Solar, the subscriber can specify an *implementation strategy* for the subscription. Because the subscriber is specifying the implementation strategy, the strategy can exploit domain-specific properties that the middleware could not know.

An implementation strategy is nothing less than a distributed algorithm for evaluating the subscription. The strategy itself is implemented with publish / subscribe amongst the brokers in the network. A strategy is kept separate from a subscription through a strategy design pattern and a factory for instantiating the subscription-strategy pair [6]. The separation provides several benefits, including the ability to:

1. parametrically substitute different implementation strategies for different environments;
2. permit non-expert programmers to specify an implementation strategy written by an expert programmer;
3. prototype an application’s behavior by writing only its subscriptions, followed later by declaratively attaching an appropriate optimizing implementation strategy;
4. reuse an implementation strategy for subscriptions that are similar in structure.

Because the subscriber is specifying the implementation strategy, the strategy can exploit domain-specific properties of the subscription that the middleware could not know. Consider our buddy proximity relationship. The first domain-specific property of note is that the given rate of location reporting is only a sampling of the actual locations. Second, if Robert is reported at location A and next at B, then he must have traveled a continuous line between those points, with time progressing smoothly over the interval. Third, there are a range of expected rates of travel: people walk at about 5 km per hour, travel in cars at 100 km per hour, and fly in planes at 800 km per hour. Fourth, the use of the proximity subscription itself, in our case, is to enable striking up a conversation.

How might these be used to improve efficiency? The distance of two people (200 km), compared against their current rate of

movement (5 km per hour) implies that they could not be proximate for 20 hours. Hence, an implementation strategy could suppress location reports at an entry node until that time has passed, either’s rate of movement changes, or either moves half the original distance between them.

3.1 Implementation Approach

Fulcrum is built on portions of Siena and Jabber [9]. To support relationship subscriptions and implementation strategies, Fulcrum employs *active subscriptions*—deployed code in the form of Java™ class files. They are integrated with Siena’s subscription and notification components. To permit interoperability with the UCSD ActiveCampus project [7], the overlay network is built on Jabber using the XML-based J-XMPP java client [10]. This allows Fulcrum to leverage the presence features of XMPP (“Robert is logged in”), enhanced with location information as detected through ActiveCampus. Support for active subscriptions and the Jabber integration required replacing Siena’s event brokers with our own. Fulcrum consists of 81 Java classes, comprising about 7,000 non-comment source lines of code.

3.2 Event Broker Architecture

Active subscriptions may take on the role of publisher, subscriber, or both. For example, our proximity subscription subscribes to the location events of two people and publishes a new event regarding proximity. In order for multiple such subscriptions to aggregate and transform events, there is a need for statefulness and multi-threading in the event broker. Our broker architecture is shown in Figure 2. There are two significant pieces as follows:

1. The primary event broker (large circle) and its associated message (event) buffer are like those from traditional CBPS systems. The primary event broker handles the normal advertisement, subscription, and publication processing. It forwards arriving events to the appropriate active subscription or another event broker. It also provides additional hooks to plug-in active subscriptions.
2. A wrapper (smaller box) provides both a sandbox and supporting interfaces for an active subscription (hexagon). It provides a




Figure 2. Event broker high-level architecture.

miniature broker environment, including a message buffer and event router (small circle).

3.3 Subscription Management

As part of our support for suppressing events at entry nodes, we make a distinction between *local* and *global* subscriptions. A global subscription is the normal subscription found in CBPS systems. We added local subscriptions to enable an active subscription to live only immediately adjacent to a data producer of interest and not have its subscription be delivered to remote pub-

lishers that might generate similar information. These are used only by implementation strategies, not by normal subscriptions. Both local and global subscriptions receive any data that arrives at a broker node. However, when a new publisher advertisement is received then only global subscriptions are propagated to broker reporting advertisement.

3.4 Active Subscription Deployment

Two things need to happen in deploying an active subscription. One, similar to Siena subscriptions, the relationship subscription needs to be deployed to every broker in the network between the subscriber and contributing publishers. Two, the implementation strategy needs to be instantiated at the appropriate brokers. Fulcrum manages deployment of an active subscription as follows.

First, the active subscription is packaged to appear as just its relationship subscription, that is, with the implementation strategy hidden inside it. Consequently, the existing subscription routing mechanisms are used. In particular, based on the registered advertisements at each event broker, the subscription follows the path for every advertisement of interest upstream until the edges of the network are reached, with the subscription being left at each intermediary node for forwarding purposes.

Second, the subscription's implementation strategy is evaluated (instantiated) at each node. For non-entry nodes, the strategy is simply discarded, and the subscription merely functions to forward events back to the subscriber. For each entry node, instantiating the strategy results in the following actions:

1. The strategy determines which events are to be subscribed to with local subscriptions. Our proximity subscription $\{(X_R - X_W)^2 + (Y_R - Y_W)^2 < 10^2\}$ implies two constituent subscriptions:

- A. $\{NAME=Robert \ \& \ X=ANY \ \& \ Y=ANY\}$
- B. $\{NAME=William \ \& \ X=ANY \ \& \ Y=ANY\}$

At entry nodes advertising Robert's location, the strategy makes a *local* subscription for A. At entry nodes advertising William's location, the strategy makes a local subscription for B. This configuration is shown in Figure 3a². Note that, after this step, the entry node for Robert (or William) will have insufficient data about William's (Robert's) location to evaluate the relationship subscription.

2. The strategy advertises new event types and makes additional *global* subscriptions that exchange data with the same strategy on




Figure 3a. Only entry nodes receive original events.

² The configuration with a third party subscriber was chosen to simplify the presentation, to demonstrate the possibility, and to avoid shortcut attempts when one buddy is also the subscriber.

the other entry nodes. The data from this sub-protocol is used in two ways. One, to *evaluate* the relationship subscription to possibly forward an event to the originating subscriber. Two, to collaboratively *schedule* the evaluation of the relationship subscription. That is, the strategy is normally evaluating ancillary subscriptions locally and only occasionally forwarding data between the entry nodes. This is how efficiencies are achieved with relational subscriptions. An example sub-protocol is shown in Figure 3b (described in full in the next subsection), and the relationship evaluation itself is shown in Figure 3c.

Note that the strategy running on each entry node does not need to know where the other strategies lie. They merely advertise their new events, and the others subscribe to them. The middleware marries them in the normal publish / subscribe fashion.

Using buddy proximity explanation, the next section gives a detailed explanation of how Fulcrum can be used to achieve efficient yet transparent context-aware publication / subscribe. Section 5 describes an experiment using Fulcrum to track buddy proximity for the ActiveCampus system, revealing the efficiencies that are possible with this approach.




Figure 3b. A sub-protocol is used to minimize updates, e.g., $O(\lg \text{ distance})$, between collaborating active subscriptions.




Figure 3c. Only a single event representing the satisfaction of the relationship is passed down to the subscriber.

4. EXAMPLE – BUDDY PROXIMITY

Evaluating relationship subscriptions at entry nodes is a necessary but not sufficient condition for the most efficient evaluation of relationship subscriptions. In particular, an efficient distributed algorithm for detection of the relationship must be implemented at the entry nodes. This section describes such an algorithm and its implementation for the proximity of two buddies. Interestingly the same approach is the most efficient solution to multiple buddy relationships, but best implemented in a single active subscription.

The essential idea is that the proximity of two distant entities cannot be achieved unless points between them are crossed: Robert cannot be co-located with William unless one or both of them first

reduces the distance separating them in half, and so on. In our *Range-Ring Strategy*, these halfway points are the only events communicated between collaborating active subscriptions at the entry nodes. All other event traffic can be suppressed.

Returning to our example of Robert and William, the idea is that when the initial active subscription is deployed, the distance between the two is computed by the strategies running at William’s node and Robert’s node. Each strategy then computes a “range ring” around its user’s position (Figure 4a), that if crossed by the




Figure 4a. Range-Ring Strategy creates range rings separated by desired proximity and suppresses all events until one buddy moves at least half the distance, such that the potential for the proximity relationship to be satisfied exists.

user, results in the strategy publishing the user’s location and computing a new range ring (Figure 4b). Because we are interested in whether *either person* crosses half the distance, any newly computed range ring is sent to the other by publishing a range ring event. When a user leaves its ring by approaching its buddy, the range rings shrink; when the user leaves its ring while receding from its buddy, the range rings grow.

To preserve consistency, a couple of things are necessary. First, the user’s new location and the new range ring are published as an aggregate event. This permits the receiver to detect what data was used in computing the range ring. Second, when the other user’s strategy receives the event, the strategy not only modifies its suppression filter to the new range, but also runs the filter against the user’s current location to see if the user has already moved beyond the new ring. If so, it publishes a new location/range-ring event to preserve the invariant that users are inside range rings.

In the case where one user moves directly towards the other, who

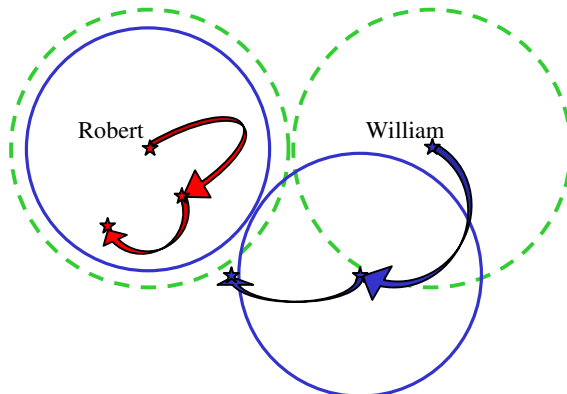


Figure 4b. Range-rings are reset when a subject moves outside its assigned area. New range-rings have solid lines.

is stationary, each successive range ring would be half the size, representing the fact that a smaller movement is required to bring

the two into proximity. For this scenario, with initial separation S and proximity bound of D , there are $\log_2(2(S-D))$ range ring events sent before the proximity event itself is sent. In a scenario where both users are moving at velocity V and remain at a constant separation S , the range-ring event *rate* is $2V/(S-D)$. The publication rate of location events does not matter in either scenario because only the significant events—those that represent leaving a range ring—are sent. Thus, the algorithm has a desirable “pay for what you use” property, and its suppression of location events has greater value as the reporting rate of the deployed positioning technology increases.

We present the following scenario for the algorithm.³ The scenario assumes that the publishers for Robert and William are already advertising their events. Figures 1 and 3a-c diagram the overlay network, publishers, and subscribers, while Figures 4a-d diagrams the algorithm’s model of the users’ behavior.

4.1 Scenario

4.1.1 Deployment:

1. To instantiate the active subscription for the proximity relationship $f(A,B,D) = \{(A.x-B.x)^2 + (A.y-B.y)^2 < D^2\}$, Subscriber 1:
 - a. binds A , B , and D to Robert, William, and 10 respectively;
 - b. identifies constituent subscription filters for Robert’s and William’s location as shown in Section 3.4;
 - c. selects the *Range-Ring Strategy*.
2. Fulcrum pushes the active subscription through the network to the event brokers adjacent to Pub_R and Pub_W respectively where it is instantiated as AS_{RW} and AS_{WR} respectively. Also, as a side effect of normal subscription deployment, the relationship subscription filter is left at each intervening node.
3. AS_{RW} publishes the range-ring event type $\{Msg=RangeRing, _name=A.name, _x=ANY, _y=ANY, range=ANY\}$, and subscribes to the event $\{Msg=RangeRing, _name=B.name, _x=ANY, _y=ANY, range=ANY\}$.⁴
4. AS_{WR} publishes the range-ring event type $\{Msg=RangeRing, _name=B.name, _x=ANY, _y=ANY, range=ANY\}$ and subscribes to the event $\{Msg=RangeRing, _name=A.name, _x=ANY, _y=ANY, range=ANY\}$.

Steps 3 and 4 are setting up a sub-protocol between Robert’s and William’s entry nodes to efficiently share each other’s location, as outlined in the following. The active subscription uses the lexicographically smaller named instance to report success to avoid a race condition that would otherwise lead to two notifications.

4.1.2 Runtime

5. Pub_R publishes Robert’s position, which is received by the adjacent entry broker and handed off to AS_{RW} only. AS_{RW} recognizes this as the first event and publishes a range ring event $\{Msg=RangeRing, _name=A.name, _x=A.x, _y=A.y,$

³ The full algorithm and its proof of correctness are beyond the scope of this paper.

⁴ Note that an underscore has been prepended to the attribute names that overlap with those of location events. This is to prevent subscribers to user location events from getting range ring events. A general solution to this problem would employ a scoping mechanism like namespaces.

- range=MAXINT}, ensuring it will not report again until it receives information about William).
6. AS_{WR} receives the range ring event from AS_{RW} and saves it.
 7. Pub_W publishes William's position, which is received by the adjacent entry broker and handed off to AS_{WR} only. AS_{WR} computes the distance between the two, and publishes the range ring event {Msg=RangeRing, _name=B.name, _x=B.x, _y=B.y, range=(distance-D)/2}.
 8. AS_{RW} receives the range ring event from AS_{WR}, saves the data, and takes the range ring as its own. At this point Robert and William both now have range rings as shown in Figure 4a.
 9. Pub_R and Pub_W continue to generate location events that are processed by AS_{RW} and AS_{WR} until such time as either Robert or William exits the assigned range ring. During this period, neither strategy is notified of the other's location, due to the event suppression imposed by the range ring strategy.
 10. William moves outside the ring first, following the arrow on the right in Figure 4a. Pub_W thusly reports William's position; AS_{WR} receives the event and recognizes it is outside the range ring. (The formula for the range-ring test is coincidentally the same as that for proximity.) Before computing a new range ring, it checks whether the proximity relationship itself has been satisfied. In this case, it has not. It then computes a new range ring based on William's current location and the location it has saved for Robert, and sends it as an event.
 11. AS_{RW} receives and saves range ring event from AS_{WR}, and takes the range ring as its own. Robert and William both now have range rings around their last reported locations, as shown in Figure 4b. Even though Robert has moved (as shown on the left in Figure 4a), because that location had been suppressed, AS_{RW} establishes the new range ring around the position it last published. This ensures that AS_{RW} and AS_{WR} have a consistent view of each other's state.
 12. Location events continue to be suppressed until William again moves outside the range ring, as shown in Figure 4b by the second arrow on the right. Robert in the meantime has followed the second arrow on the left. Pub_W publishes William's position; AS_{WR} receives the event and recognizes the location as outside of the range ring. As in step 10, a new range ring is computed and published as an event, in effect creating the smaller range rings shown in Figure 4c.




Figure 4c. Updates use most recently reported position of each subject, resulting in concentric rings for slow movers.

13. AS_{RW} receives the event data with the new range ring. This time, it recognizes that Robert's most recent location is outside the smaller range ring that AS_{WR} created for it. If taken as-is, AS_{RW} would never see Robert leave the ring. Consequently, it computes a new range ring and sends it as an event, creating the small range rings in Figure 4d.




Figure 4d. Range-rings are reset when a subject is discovered outside a newly assigned area.

14. This pattern repeats until the proximity relationship is satisfied, at which time AS_{RW} (the smaller named applet) publishes the proximity event, which is routed back to Sub1.

4.2 Discussion

In summary, active subscriptions replace numerous user location events with a few self-regulating range-ring events, thus substantially reducing the event traffic for computing buddy proximity.

The use of range rings to indicate when to next publish positional information is only one possible implementation strategy. We could have identified the perpendicular line between the two positions and then created a pair of parallel lines 5 meters on either side of the center as shown in Figure 5. Then the active subscrip-




Figure 5. Parallel-line strategy bisects space

tions would monitor for events revealing that their subject has crossed the line. Again the positions and distances would re-sync. If one of the actors was known to move faster than the other, the parallel lines or the range rings could be skewed in position or size to account for the additional context.

If a subscriber wishes to use the same concept of proximity, but across different dimensions, such as temporal or financial, then perhaps an existing strategy could be used. Consequently, a library of strategies could be a useful resource in an open publish / subscribe system.

5. EXPERIMENTAL EVALUATION

In theory, the Fulcrum open implementation approach permits users to apply domain knowledge and evaluate context-aware relationships at entry nodes, while suppressing useless events. However, the efficiencies achieved by the Fulcrum approach are

dependent on properties of the supported applications, the way people use them, the state of the underlying technology, and the structure of the event broker network.

This section describes an experiment of Fulcrum supporting a federated architecture for ActiveCampus and compares those results to the basic and enhanced CBPS approaches. We perform this evaluation in two parts. First, how much suppression of original events does Fulcrum achieve? Second, what overall efficiencies are achieved?

5.1 Experiment

ActiveCampus user positions are automatically generated via triangulation based on 802.11 signal strength and automatically reported to the system. On ActiveCampus’s “buddy page”, the user display shows all buddies with an indicator of *nearby* or *far*, mirroring the proximity relationship that we have used in our examples. Consequently, we took a week’s worth of data from ActiveCampus and “replayed” it through Fulcrum, using a conservative broker configuration.

5.1.1 Setup

The ActiveCampus test data is comprised as follows:

- 643 users (anonymized for privacy).
- 2,165 buddy relationships
- 604,800 seconds = 168 hours = 1 week
- 1-6 events per logged-in user per minute
- 360,067 location reports

The overlay network was configured with 5 event broker nodes in a crossbar configuration with one internal “routing” node (Figure 6). Such a configuration might be used geographic coverage (e.g., one node per city plus the crossbar) or to mediate different com-




Figure 6. Crossbar event broker overlay network.

mercial providers (e.g., AIM, MSN, ICQ, and Yahoo connected through a crossbar). Different configurations would change the number of hops that an event has to travel. This configuration has a maximum hop count of 2, minimizing the penalty for an event not being suppressed.

The 643 users were randomly distributed across the 4 edge brokers, resulting in 149, 158, 165, and 171 users at each respective node. Each user was configured as a publisher of its location as well as a subscriber to the proximity of each buddy.

This configuration resulted in the distribution of proximity subscriptions as shown in Table 1. Users attached to node 0 have proximity relationships with the number buddies on nodes 0 to 3 as 125, 155, 165, and 112 respectively. This resulted in 2,165 active subscriptions being deployed. Each was instantiated on two nodes, yielding 4330 instances.

Nodes	0	1	2	3
0	125	155	165	112
1	131	140	167	115
2	125	171	224	126
3	85	110	122	92

Table 1. Proximity relationship distribution.

5.1.2 Results

Of the original 360,067 location reports, 11,955 were associated with users not participating in a buddy relationship, leaving 348,112 location reports that were subject to the proximity subscription.

Many users were not logged in to ActiveCampus at the same time as their buddies during the one-week period. Of the 2,165 active subscriptions, then, only 1,028 received data from both buddies. This means, effectively, that the 348,112 location reports were concentrated on half of the users. The range ring strategy resulted in 57,348 published range-ring events and 4,626 proximity relationship events (i.e., reports of a buddy moving into proximity) – 16.5% and 1.3% of the original event count respectively, for a cumulative reduction of 82%. The breakdown of event reduction for each user’s buddy is shown in Figure 7.




Figure 7. Range-ring events generated per original event. A line is overlaid to show an $O(\lg_2 \text{ movement})$ expectation.

Comparing the number of original events to the number resulting from the active subscriptions yields an average event savings of 82% (a factor of 5.6). The average savings factor does not tell the whole story, however. Due to the centralized architecture of ActiveCampus at the time of data capture, user events were rate limited to 1 to 6 events per minute. In the wild, it is not unreasonable to expect 1-second (standard GPS reporting rate on naval ships) or better update rates – a 10-times increase. Because the number of range-ring events is based on distance, not the number of reports, the average event savings factor would then be 56. Also, user activity plays a significant role in event reduction as shown by a few sample data points displayed in Table 2.

User Activity	# orig events	# range ring events
Stationary	22,657	6
Stationary	11,743	2
Mobile	1,008	127
Mobile	4,746	145

Table 2. User activity affects event reduction.

In many cases we find that users are relatively stationary. Their usage model is to go somewhere, turn their device on, remain in a

confined space for a time, and finally to turn off their device and repeat the cycle. Thus, almost all events are suppressed at the source. Conversely, when people are in motion, there is usually much less suppression because buddies frequently exit their assigned range rings yet seldom walk directly toward each other.

5.2 Analysis

Although we have achieved an 82% reduction in events, any efficiency gained must be evaluated from a systems perspective. In particular, the question is how much processing the whole system has to perform for all the events communicated.

Due to difficulties in running live Siena, Gryphon, and Solar CBPSs on the ActiveCampus data, we measured system efficiencies in terms of analytically accessible properties. The core costs are the processing of an event at an event broker and any subsequent forwarding of the event through the broker network. We use the number of *hops per original event*, or *event hops*, as our common unit of measure to evaluate the efficacy of Fulcrum. We discuss the cost per hop in the next section.

To simplify the calculation of effectiveness we assume a uniform distribution of publishers, subscribers, and active subscriptions across all edge nodes. To verify that this simplification would not bias our results, we compared Fulcrum’s empirical aggregate hop count based on the random distribution to the analytically derived one. The empirical aggregate count based on the configuration in the previous subsection is 94,334 hops. Analytically, it computes as 94,260, less than 0.1% smaller. We have confidence, then, that the analysis here is accurate.

Basic CBPS like Siena requires the subscribers to evaluate all relationships. Consequently, $\frac{1}{4}$ of the events coming into entry nodes and are immediately returned to subscribers attached to the same node, resulting in 1 event hop apiece (we only count processing and output, not the input). $\frac{3}{4}$ of the events will come into an entry node and get passed to the center node, on to one of the remaining three edge nodes, and finally to the subscriber, yielding 3 event hops. The subscriber itself evaluates the relationship requiring 0 event hops. On average this is

$$\frac{1}{4} * 1 + \frac{3}{4} * (3 + 0.0133 * 0) = 2.5 \text{ event hops.}$$

Aggregation enhanced CBPS like Gryphon or Solar allows relationships to be computed at the first common node (1stCN). For the given broker configuration, $\frac{1}{4}$ of the events come into an entry node that also acts as a common node and the proximity relationship “success” event can be generated and sent to the subscriber in 1 event hop. $\frac{3}{4}$ of the events come in through their entry nodes and are passed to the center node that is consequently the first common node, yielding 1 event hop. Success events generated at the center node must be communicated with the subscriber for 2 more event hops. On average, this is

$$\frac{1}{4} * (0.0133 * 1) + \frac{3}{4} * (1 + 0.0133 * 2) = 0.773 \text{ event hops.}$$

Open-implementation CAPS like Fulcrum’s reduces original events into potentially significant events at the entry nodes. In this case, $\frac{1}{4}$ of the events come into an entry node that possesses both parts of an active subscription, yielding 1-hop success events. $\frac{3}{4}$ of the events come into an entry node hosting an active subscription that collaborates with a distant entry node. Empirical data from the previous sub-section shows that 6.07 original events will become one collaboration event (range ring). These events

travel through the center node and out to an associated active subscription. This yields 0.165 event hops for such events. Success events are then passed to the subscriber, which half the time will be local (1 hop) and the other half will be remote (3 hops) for an average of 2 hops. The net event hops per original event are

$$\frac{1}{4} * (0.0133 * 1) + \frac{3}{4} * (0.165 * 2 + 0.0133 * 2) = 0.27 \text{ event hops.}$$

When location sensor reporting is increased to 1 report per second, a 10-times increase, then basic CBPS will experience a 10x increase in location events at 2.5 event hops apiece. Gryphon and Solar would experience a 7.5 times increase in location events to the 1stCN, and maintain their 0.773 event-hop count. In Fulcrum, because forwarded events are driven by user behavior, not event rate, the overlay network sees no increase in traffic, which yields a corresponding 10-times effective suppression, resulting in 0.027 event hops per original event. Solar could achieve similar reductions by adding a transformation operator that rate-limits original location events, but would sacrifice the added accuracy of the increased reporting rate unless the rate-limiter made inferences over all the events.

6. DISCUSSION

Filtering is the lifeblood of CBPS. Our ability to reduce event traffic by filtering relational events at entry nodes is our key efficiency measure. We now discuss some of the tradeoffs that arise in achieving this reduction.

One concern is the large number of subscriptions that might be deployed as Java code. For our experiment, there were 2,165 relationship subscriptions for 643 users with 3.37 average buddies. It should be noted, first of all, that there is intrinsic sharing. The Java class loader only loads unique classes, so only the first such subscription at a node results in transmitting code and configuring the class. Second, open implementation allows the subscriber to achieve higher-levels of sharing by plugging in a new strategy. For example, a user could attach a strategy that tracks multiple buddies for a user, resulting in a sharing of range rings and event subscriptions behind the scenes.

Another question is whether subscribers could, with an existing CBPS like Solar or Gryphon, write collaborative subscriptions like those realized through our open implementation approach. Indeed, for our example, a subscriber could put out subscriptions for range ring events on each publisher. The subscriber itself would have to act as the clearinghouse that receives the events. Upon doing so, it would create new range ring subscriptions and retract the old ones. There are three problems. For one, this added level of indirection would increase the hop count unless the subscriber was attached to the 1stCN. Two, there is not a clear separation between the relational property of interest and its efficient implementation, increasing complexity and reducing the opportunities for reuse. Three, there is a race condition, at least with this straightforward adaptation. Our algorithm avoids the race condition by *updating* the existing subscription on the other entry node, and before doing so checks that the other user has not already moved outside the new range ring. This kind of atomic check is not possible in normal CBPS when placing a new subscription; a subscription filter simply waits for the arrival of the next event.

Although hop counts are reduced by our approach, event processing is slower than in a basic CBPS like Siena. In providing the ability to perform attribute-to-attribute comparisons, efficient

structuring of subscription evaluations by event brokers (e.g., exploiting overlaps, like $X > 10$ is subsumed by $X > 5$) is hampered. Solar and Gryphon share this problem. To mitigate this, we take a cue from database query optimization and reorganize subscriptions to evaluate the “normal” part of a subscription first, and only perform the attribute-attribute comparisons if it succeeds. More aggressive optimization is a subject of future work. It should be noted however, that reducing the hop count through open implementation is scalable with respect to increasing sensor reporting rate, whereas as basic CBPS is not.

The Java code of an active subscription could be inefficient, buggy, or malicious. This problem is mitigated by sandboxing the computation in a separate thread and using mechanisms akin to Jabber’s *karma*, which reduces the resources available to “greedy” subscriptions. Still, the question remains why a publisher would run code for a subscriber. In the deer proximity example, the deer do not get a say. For buddy proximity, they have explicitly agreed to trust each other through a mutual buddying subscription.

7. CONCLUSION

The commoditization of networked sensors is fueling the emergence of internet-based context-aware applications. Content-based publish / subscribe (CBPS) systems are a natural substrate for supporting context-aware application development because they provide for separation of publishers and subscribers, efficient event distribution, extensibility, and scalability. Yet, the CBPS notion of subscribing to a publisher’s content does not capture the data relationships across publishers that drive context-awareness. Recent advances in aggregation-enhanced CBPS permit subscribing to these relationships, but the relationships are computed at intervening nodes in the CBPS middleware, resulting in potentially large amounts of raw data being routed to those nodes.

The problem is that the separation of concerns afforded by the middleware precludes publishers’ event brokers from collaborating with subscribers and each other to implement algorithmically efficient application-specific context-aware inferences. It is possible to dramatically reduce the number of events processed and forwarded through the middleware by using open implementation to permit subscribers to enhance the middleware with the ability to deploy a distributed algorithm to where events enter the system. Open implementation provides for separation of subscriptions and implementation strategies, allowing for separate, modular development and reuse of relationship subscriptions and implementation strategies.

Fulcrum supports open-implementation context-aware publish / subscribe (CAPS). It employs *active subscriptions* in the form of Java applets. They are first-class entities that both subscribe to events and publish new events, enabling event-based collaboration amongst copies of the strategy deployed where raw sensor events enter the system.

We evaluated the Fulcrum approach by implementing a buddy-proximity relationship and accompanying strategy. The resulting strategy’s performance is tied to the amount of user movement, not to the number of events reported. We then federated the ActiveCampus context-aware infrastructure and replayed a week’s worth of its user data. For a conservative broker network configuration and current sensor technology, we found an 82% net event reduction and a 64% reduction in event hops compared to aggrega-

tion-enhanced CBPS, which uses similar underlying technology. These factors increase proportionally with increased reporting rates, naturally insulating against advances in sensor technology. The data curve-fits to $O(|\lg_2|\text{movement}|)$ expected event processing.

8. REFERENCES

- [1] Carzaniga, A., Architectures for an Event Notification Service Scalable to Wide-area Networks, PhD. Thesis, Computer Science, University of Colorado, Boulder, 1998.
- [2] Carzaniga, A., Rosenblum, D. S., and Wolf, A. L., Achieving Scalability and Expressiveness in an Internet-Scale Event Notification Service, *19th ACM Symposium on Principles of Distributed Computing*, pp 219-227, 2000.
- [3] Chen, G. and Kotz, D., Solar: Toward a Flexible and Scalable Data-Fusion Infrastructure for Ubiquitous Computing, *UbiComp 2001 UbiTools Workshop*, 2001.
- [4] Chen, G. and Kotz, D., Solar: A Pervasive-Computing Infrastructure for Context-Aware Mobile Applications, Dartmouth Technical Report, TR2002-421, 2002.
- [5] Cohen, N., Purakayastha, A., Turek, J., Wong, L., Yeh, D., Challenges in Flexible Aggregation of Pervasive Data, IBM T.J. Watson Research Center, Tech Report RC21942, 2000.
- [6] Gamma, E., Helm, R., Johnson, R., Vlissides, J., *Design Patterns: Elements of Reusable Object-Oriented Software* Reading, MA, Addison-Wesley, 1995.
- [7] Griswold, W., Boyer, R., Brown, S., and Truong, T., A Component Architecture for an Extensible, Highly Integrated Context-Aware Computing Infrastructure, *2003 International Conference on Software Engineering*, May 2003.
- [8] Griswold, W., Shanahan, P., Brown, S., Boyer, R., Ratto, M., Shapiro, R., and Truong, T., ActiveCampus – Experiments in Community-Oriented Ubiquitous Computing, Technical Report CS2003-0750, Computer Science and Engineering, UC San Diego, June 2003.
- [9] Jabber Software Foundation. 2004. Jabber: Open Instant Messaging and a Whole Lot More, Powered by XMPP. [Online] Available: <http://www.jabber.org/> [2/22/04]
- [10] JabberStudio. 2003. J-XMPP Project Information, [Online] Available: <http://www.jabberstudio.org/projects/j-xmpp/project/view.php> [2/22/04]
- [11] Jin, Y. and Strom, R., Relational Subscription Middleware for Internet-Scale Publish-Subscribe, *Proceedings Second International Workshop on Distributed Event-Based Systems (DEBS '03)*, 2003.
- [12] Kiczales, G., Beyond the Black Box: Open Implementation, *IEEE Software*, January 1996.
- [13] Kiczales, G., Lamping, J., Lopes, C.V., Maeda, C., Mendhekar, A., Murphy, G., Open Implementation Design Guidelines, *1997 International Conference on Software Engineering (ICSE)*, May 1997.
- [14] Tarasewich, P., Toward a Comprehensive Model of Context for Mobile and Wireless Computing, *Proceedings of America’s Conference on Information Systems (AMCIS) 2003*, 2003.