

IBM Research Report

Mobile Devices Meet Customer Loyalty First Of A Kind (FOAK) Project Summary

Alan Cole, Scott McFaddin, Chandra Narayanaswami, Alpana Tiwari
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598



Mobile Devices Meet Customer Loyalty First Of A Kind (FOAK) Project Summary

Alan Cole, Scott McFaddin, Chandra Narayanaswami, Alpana Tiwari
IBM TJ Watson Research Center, Hawthorne, NY

Abstract: *IBM Research recently completed a First-Of-A-Kind (FOAK) engagement and pilot program with a major grocery retailer that integrates mobile device usage with the retailer's loyalty program. This project, entitled "Mobiloyalty", enables loyalty program customers to use their mobile devices to receive in-store promotions when they visit a retail store, with discounts applied automatically at the point of sale. Mobiloyalty builds upon a base technology, called Celadon, which integrates mobile devices with services and information offered in public spaces.*

1: Introduction. Mobile devices are increasingly being recognized as important channels by which retailers may drive transactions and manage relationships with their customers. Of paramount importance for retailers is the ability to offer services, information, and transactions to customers on their mobile devices *while they are inside a retail store*. To address this need, IBM Research and IBM's Ubiquitous Computing Laboratory completed a base technology in 2008, called Celadon, for integrating customer-owned mobile devices with transactional and informational services in general public spaces. Celadon is a software framework which decomposes mobile commerce scenarios into a battery of core elements (e.g. location sensing) and exposes those elements as REST based services. Celadon also provides tools and utilities for integrating those services with mobile devices and retailer systems. Celadon is itself domain independent, but supports the addition of domain-specific frameworks, such as retail.

As a real world test of this technology, IBM partnered in 2008 with a major U.S. retailer in a publicly piloted solution, called Mobiloyalty, integrating consumer owned devices with the in-store shopping experience, enabling customized, on-the-spot, in-store promotions and offers. The Celadon/Mobiloyalty solution unified the retailer's loyalty program systems, with the loyalty program customers' mobile devices enabling customers to receive price promotions that may be redeemed at a point of sale when they visit a pilot store. Although the loyalty program was the focus of our pilot, the design of the Mobiloyalty system, described below, can accommodate customers who are not loyalty program members. This solution was piloted between June and September of 2008 at a single store selected by the retailer. Usage and traffic data collected during the pilot are provided below.

This paper is organized as follows: Section 2 describes the operation of the solution and analyzes its three major components. Section 3 describes the design principles and architecture employed to operate the scenario. Section 4 contains a set of observations and lessons learned from the operation and implementation of the scenario. Section 5 contains the numerical and anecdotal measurements of the system (usage rates etc.), makes suggestions for future research and development, and presents our overall conclusions.

2: Solution Scenario.

The Mobiloyalty solution was designed to engage the user throughout the lifespan of the program via both user-initiated interactions as well as system-initiated interactions. The solution operated both in-store and out-of-store. There are three parts to the scenario, as illustrated in Figure 1.

2.1 Enrollment.

The enrollment phase incorporates participants into the program. The goal of the enrollment phase is to collect user information which may be subsequently brokered to the retailer for purposes such as recognition of users in stores, and enabling the initiation and offering of services based upon the brokered information.

The enrollment phase establishes an association between the user's identity, their mobile device information (such as phone number), their loyalty program identifier, and individual shopping preferences (for example, how information should be channeled to them and under what circumstances). The enrollment process also establishes correlations between the user's identity and physical identifying information associated with the user or their mobile device.

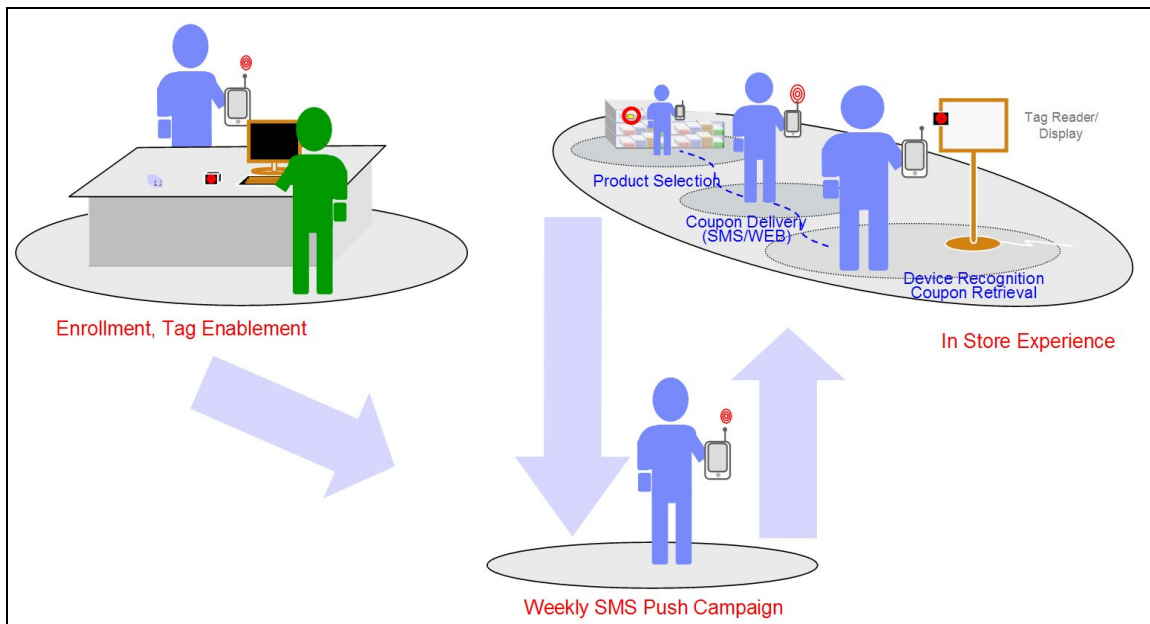


Figure 1. Illustrating the three phases of the Celadon/Mobiloyalty in-store mobile commerce solution.

In the case of our pilot we used a *near field communications (NFC)* strategy for customer recognition. A small electronic tag, approximately the size of a U.S. quarter dollar coin, was attached to each user's mobile phone for physical identification. The purpose of the electronic tag is also explained to users during the enrollment phase. Each tag has a unique identifier that allows the phone to be recognized, when placed near public reader points in the store. As the tag has a short range (less than 10cm), this approach does not allow the phone to be recognized involuntarily without the user's knowledge. More recent smart phone models are natively capable of near field communications -- the tag and circuitry are built in -- the approach of this paper applies without the need for physical outboard tags.

2.2 In-Store Experience.

Upon entering the pilot retail store, the customer "taps-in" by waving their device near a reader station deployed at the store entrance. This physical recognition causes the user's identity to be de-referenced from the enrollment information collected in the enrollment phase, and the location

of the user to be reported to the Celadon/Mobiloyalty system. The system accesses the customer profile, preferences, and loyalty program account, and determines an appropriate service for responding to the detection of the customer at the given location, in this case, store entry. The responding service contacts the retailer's promotion system, which computes a set of customized promotions on select products for that customer. The promotions system can be supplied with and leverage rich information about the customer's context, e.g. their location context, which includes the specific store location and even location within the store of the reader point, enabling them to make meaningful promotional offers.

Upon determination of a set of promotions applicable to the context of the user's store visit, the responding service inserts those promotions into a promotions repository (a "digital wallet") associated with the user and transmits a description of the promotions to the customer via an SMS text message. The SMS message also provides a URL which references a web page compatible with their mobile device, customized to the user, containing a more detailed version of the promotions data plus detailed product information, price discount information, and images further describing the promotions. The information directs the customer to select products that the retailer wishes to highlight or to promote cross-sell and up-sell purchase patterns. The information displayed on the device may be used by the customer during their shopping experience to locate and select offered items that they are interested in. Finally, the price discounts are applied automatically at the point of sale, when the featured products are found in the user's basket and correlated to the pilot program.

The promotions offered in Mobiloyalty have a limited lifespan that operates on cycles which expire at the end of the current week, this may be generalized to arbitrary life spans. Additionally the promotions extended to a user are always included in their digital wallet, and are always considered enabled. In more general schemes of managing digital wallets, promotions could be subject to more flexible *organizational schemes* in which they are selectively accepted or rejected, classified, tagged and/or assigned to folder in some intuitive hierarchy. E.g. the digital wallet might have a folder representing a "coupon inbox" in which offers are staged. Further, promotions could be subject to *status management schemes* in which they are selectively enabled or disabled. Both organizational and status management schemes may be operated by the user manually or by some proxy agent acting on behalf of the user, the digital equivalent of selectively clipping coupons from a physical medium.

2.3 Opt-in Reminder Campaign.

A few weeks into the pilot, we introduced a program in which messages containing reminders about the week's special promotions were sent to the mobile phones of participants. These reminders contained a synopsis of the promotions. The objective of the reminder campaign was to provide a benefit both to the retailer and to the user: for the retailer, the reminders stimulated traffic for the program, and for the user, the messages provided a reminder of an activity that they considered useful.

The transmission of the reminders was keyed to a particular point in time that would have the highest benefit for both players. Analysis of the tap-in data up to this point showed that some shoppers exhibited very consistent behavior patterns, typically shopping at the same time of day (e.g., during the lunch hour, after work) and/or the same day of the week (e.g., Wednesday after work or Saturday morning). The Mobiloyalty reminders were sent to these consistent shoppers a few hours in advance of their usual shopping time; others got the reminders at a standard, mid-week, time.

We found this simple mechanism very effective in keeping user traffic data high. However, future research is needed into the full body of motives and tolerances of this type of push message. In a

general model this mechanism should be viewed as a competitive asset for reaching the customer, keeping in mind that too many reminders start to look like spam.

3. Technical Challenges and Solutions.

There were several key technical challenges that had to be overcome to enable Mobiloyalty as a real world solution. First, there was a need for an inexpensive and practical means of recognizing the customer inside the store, and securely brokering the information to the retailer. Second, there was a need for appropriate business processes, which operate independently from the means of recognition, to be matched to the customer, provisioned and initiated at appropriate points in time. Third, there was a need for content to be channeled to the user's device according to its capabilities, the system should be usable by any mobile device user, not just those with advanced phones having sophisticated capabilities. Fourth, the overall system must be easily configured, managed and monitored by the retailer.

Our overall approach was to build the Mobiloyalty system upon the design principles and operating infrastructure of the Celadon technology, described previously [1]. Using Celadon,, mobile commerce scenarios are analyzed and constructed in a component-oriented fashion. The components are expressed as REST based services which model a particular function as a collection of nouns (data definitions) against which a common set of verbs apply. For example, one core component is a location reporting component in which a noun captures the combination of a customer identity and a store location. Common HTTP-style verbs such as GET, POST, PUT and DELETE then model customer lookup, customer arrival, customer movement and customer departure, respectively. As the REST based services are agnostic about the type of component reporting (noun,verb) pairs, this affords an elegant separation of logical and physical functions into pluggable elements. Consequently most of the elements of the Mobiloyalty design may be substituted with different implementations, yielding various alternates for things like customer recognition, responses, etc. As such, Mobiloyalty can serve as a conceptual template for applying the Celadon design [1] to a wide array of uses centering around on-the-spot in-store interactions with retail users.

3.1 Decoupling Recognition from Response.

A very important design principle for the application was to decouple the business logic from particular methods of reporting/detecting location and other contextual information.

As noted above, customer recognition was accomplished by attaching inexpensive electronic tags to the mobile devices of users when they registered in the program, deploying tag readers in the physical store at desirable interaction points, and monitoring those readers for the arrival of tags. The tag used in this project is an ISO 14443 compliant, wafer-shaped, passive (no battery), RFID transponder, approximately the size of a U.S. quarter-dollar coin. Using this or any similar short-range RFID or NFC (Near Field Communication) technology provides an inexpensive and practical approach for device recognition in mobile device markets where in-board RFID or NFC devices are not prevalent. This approach yielded a simple and unobtrusive way for the customer to signal their presence in the store with a simple gesture, without the extra steps of having to interact with a mobile device application. Although elegant, our design was not hard-wired to this particular form of recognition.

We used the Celadon architectural approach in which recognition is separated from response. As noted, this was accomplished with a set of core REST services that manage each user's context (e.g. location, identity, preference). Contextual information is initially reported by recognizer modules. The reporting modules can run in the infrastructure, on the device, or possibly by third party services. Responding modules monitor those REST services based on filters and patterns which

enable them to select which types of events to receive (e.g. to respond whenever any user is added whose context points to a specific location such as "Retailer1.Store1.WestEntrance"). Responding modules can also write to the user's context, and may also be chained together.

For example consider the following chain of reporting/responding modules (we have simplified the user context to a list of name-value pairs):

Step 1: RFID tag reader module A detects a specific tag or device ID (coupled with data records read from the tag), adding a new user context record to the user context service:

```
Source: Reader Module A
DeviceID: Tag B
DeviceData: <null>
Timestamp: 7:02:22.351 PM PDT
```

Step 2: Responder module B monitors the service for the addition of any new user context records and looks up the user identity, loyalty card number and mobile phone number associated with the tag ID and data, and also the fact that the reader is posted at a known logical location, updating the user context record as follows:

```
Source: Responder Module B
DeviceID: Tag B
DeviceData: <null>
Timestamp: 7:02:22.688 PM PDT
MemberID: John Smith
ZoneID: Store1951
CellID: Entrance Foyer
MobilePhoneNumber: 925-555-1212
LoyaltyNumber: 123456789
State: Arrival
```

Step 3: Responder module C monitors the user context service for any user context records in the "Arrival" state and initiates a business process which decides to extend on-the-spot offers to the user. Responder module C also upgrades the user's context to a state of "OfferComputed" and appends the user's context with a list of offers (the below is simply a summary – the actual offer data has much more detail).

```
Source: Responder Module C
DeviceID: Tag B
DeviceData: <null>
Timestamp: 7:02:23.012 PM PDT
MemberID: John Smith
ZoneID: Store1951
CellID: Entrance Foyer
MobilePhoneNumber: 925-555-1212
LoyaltyNumber: 123456789
State: NewOffers
Offers: {"2.50 off Acme Paper Towels", "50 percent off any wine purchase"}
```

Note the chaining, starting with simpler modules which upgrade the user's context information from one form to another, for example, converting physical location data to logical location data. Other response modules handle the task of dereferencing the user's identity from device or tag information, brokering their higher level profile (e.g. their shopping preferences) into their local user context. Higher level response modules can respond once a sufficient amount of user context is collected to implement a business process. In the Mobiloyalty scenario, the business process is the

promotions retrieval described in Section 2.2.

There are numerous alternate schemes for location reporting modules which could be applied to retail scenarios such as Mobiloyalty. These schemes differ in various ways and in general location information can be reported in physical terms (e.g. {lat,long,alt}) or logical terms (e.g. {zone, cell, location} – see [1]) and cross translated. Several examples apply:

- Check-in: These schemes are essentially location reporting modules that report a logical location based upon a voluntary user report. It would be straightforward to integrate the Mobiloyalty scenario with various check-in reporting systems.
- Scanning: In other schemes the user may scan a physical pattern such as a two dimensional QR code, which contains encoded information that is translated into a logical or physical location report. This can be done either by a module which maps data encoded into the pattern into merchant location info, or by directly encoding the merchant location info into the pattern
- GPS-based geo-fencing: These types of services can be used as location reporting modules, in most cases, handling intensive low level physical reporting streams, refining them into lower intensity logical or physical location reports against known locations or regions. The refinements are based upon knowledge of particular geographical constructs such as store locations or regions (fences) around them, monitoring user location in some continuous, semi-continuous fashion or user-initiated fashion, and reporting changes in location or boundary crossings as updates to the user's logical location context.

Other schemes work in passive ways, in which no user action is required – for example, monitoring a store WiFi network for the presence of the user's device, and reporting location information when new devices are discovered that correlate to registered users.

3.2 Active Management of Shopper State.

Another principle from the core Celadon design that we took advantage of was the management of user state. The architecture described in [1] provides for a general state machine which operates in conjunction with the core Celadon data services and which provides the core state information that drives the state transition logic.

We used a simple version of this capability in which user transient state is represented in a single scalar value (the "State" field seen in section 3.1) in the user's context records. In our solution, users are cycled linearly through a sequence of fixed states as they progress through the store visit. The initial state ("Arrival") reflects simply that the user has been detected, the next state ("NewOffers") reflects that offers have been computed by the retailer's promotion engine, and subsequent states reflect various stages in the issuance of those offers.

In our application, the state transitions are implemented directly in Java code, as a set of listeners assigned to the core user context services. Each state transition is represented by a Java block which triggers when a particular state is entered, performs a computation block (e.g., triggers the offers business logic, then updates the state to the next value.

User state and the associated details of their overall context are made visible to all of the components of the system via the REST services, allowing the solution logic to be distributed among different machines and locations. In particular we operated the solution at two locations: one in the pilot store itself, and the other in the "cloud", via a REST server located at our Research facility. The in-store components were triggered by states early in the process – this is due to the need for tight coupling with the recognition process. Other components become involved later in the process, for

example, the mobile web presence operates once the promotions are issued later in the process.

3.3 Integration with Marketing and Promotions Engines.

A core aspect of the system was the linkage between business logic modules, which make high level customer relationship decisions, with the user and store contextual information managed by the system. The system serves as both a triggering mechanism for business logic modules (i.e. when to initiate them) and as a brokering mechanism that channels user and store context to those modules (what they should do, what context do they need).

In our application the retailer provided a business logic module which integrated with their existing marketing and promotions engine. There are a wide variety of such engines available to retailers, and generally they are capable of making customized decisions about individual shopping patterns and potentially useful offers or other business relationships that may be suggested to users. Our system added a key element to the equation: the ability to trigger and provide context for computing on-the-spot marketing opportunities and then channeling those opportunities to the user via the mobile device.

In other cases the same pattern could be applied to other business process modules operated by a retailer or site operator. For example, customer processes such as package pickup could be integrated in a similar way, making use of customer time, location and other contextual information to make the experience more efficient and improve customer relationships and increase sales.

The triggering and context brokering logic for the in-store scenario (2.2) was reasonably simple – there was a single point of interaction, correlated with a single business function (offer delivery) and the interaction occurred at a known time (the time of customer recognition via the tap-in). For other interactions, more highly customized business logic must be applied. For example, the reminder campaign (2.3) was keyed to the application of a business function (generation and push of a reminder message) at a specific opportune time. In the case of the pilot the opportune time was a few hours before a mean shopping time, customized to each shopper. In this case, a background daemon process was used (similar to a Unix “cron” or a Window “at” process). Instead of relying upon an instantaneous trigger such as that used by responder module C in Section 3.1, the daemon process would begin its business process at pre-computed time points which were customized per-user based on analytical reasoning (see below).

Engineering note: for the sake of robustness we also backed up the business process modules with fixed modules that operated from a data “sandbox”. In cases where the more advanced business logic modules, such as the promotions engine, was unavailable, the system rolled over to a fixed store of offers which could be issued to all users, without the need for individual customizations.

3.4 Integration with Analytics Engines.

One of the chief goals of the project was to collect usage data and to analyze that usage data at the end of the pilot. However, mid way through our pilot, it became clear that there was a great deal of value to be derived from the usage data as it was being collected. In particular, the question arose as to how to dynamically use information about usage patterns to improve the efficacy of the user experience. Could the system keep itself in the front of the user’s mind as the pilot transitioned from an initial novelty to something that could become part of everyday mobile commerce?

Thus, we tailored the Mobiloyalty solution to (1) collect and periodically report summaries of the user’s context (location of use, time of use, activity) as users engaged the system, and (2) channel those reports to analytical modules, and (3) correlate business functions with the output of those analytical modules.

User context was collected in the form of activity/response records, which described time and location plus the type of a particular activity. An example of an {activity,response} pair was {user tap-in, offers issued} as seen in step 3 of section 3.1. The activity/response records were collected by an activity data service, a REST based service that operates along the same design as the other core Celadon services.

Analytical modules could then subscribe to the activity data service, filtering and receiving reports which were customized to particular types of activity, response, or user context. In the case of the pilot, Celadon/Mobiloyalty provided two analytical modules that were capable of reviewing shopper activity and making deductions about shopping patterns. The first computed a mean time of week that the shopper tends to tap-in and a standard deviation around that point in time. This analytical module is effective in identifying customers who have a regular time of week that they shop for groceries (e.g., Saturday morning) – such customers are identified by a small deviation interval. The second analytical module computed a mean time of day that the shopper tends to tap-in, coupled with a standard deviation. This module is effective in identifying customers who shop at a regular time of day (e.g. they visit the grocery store during their lunch hour or on the way home from work).

Each run of the analytical module was collected into an overall report, which contained, for each user, an analysis of their shopping pattern as described above. To complete the loop, Mobiloyalty used the analytics reports to schedule the triggers for the weekly reminder campaign described in section 2.3. The reminder campaign used a default time at which reminder messages were transmitted to all users. Based on the analytical report, if a strong time of week or day pattern was observed, the reminder time was adjusted based on their anticipated shopping time (sending the reminder message a few hours in advance of the estimated time). In the future we envision that this information could also be coupled with wish lists and similar information to provide even greater anticipation of the customer activity.

Figures 2 and 3 show Celadon/Mobiloyalty administrative consoles for managing the analytical capabilities of the system and for associating those capabilities with business processes. In Figure 2, a multi user view is seen. The graphical columns carry time of week and day histograms for tap-ins. Figure 3 illustrates a single user view summarizing the weekly and daily traffic via cyclical histograms. In both diagrams the orange dots identify the mean tap-in times and the purple regions indicate the deviation around the mean time. A smaller deviation means a regular shopping pattern has been detected. Note that some customers may have a regular weekly pattern but an irregular daily pattern and vice versa. The customer seen in Figure 3, for example, has no discernable weekly pattern but a highly regular daily pattern, with a mean visit time in the 5-6 PM hour, and with a high reliability. This customer most likely visits the grocery store after work.

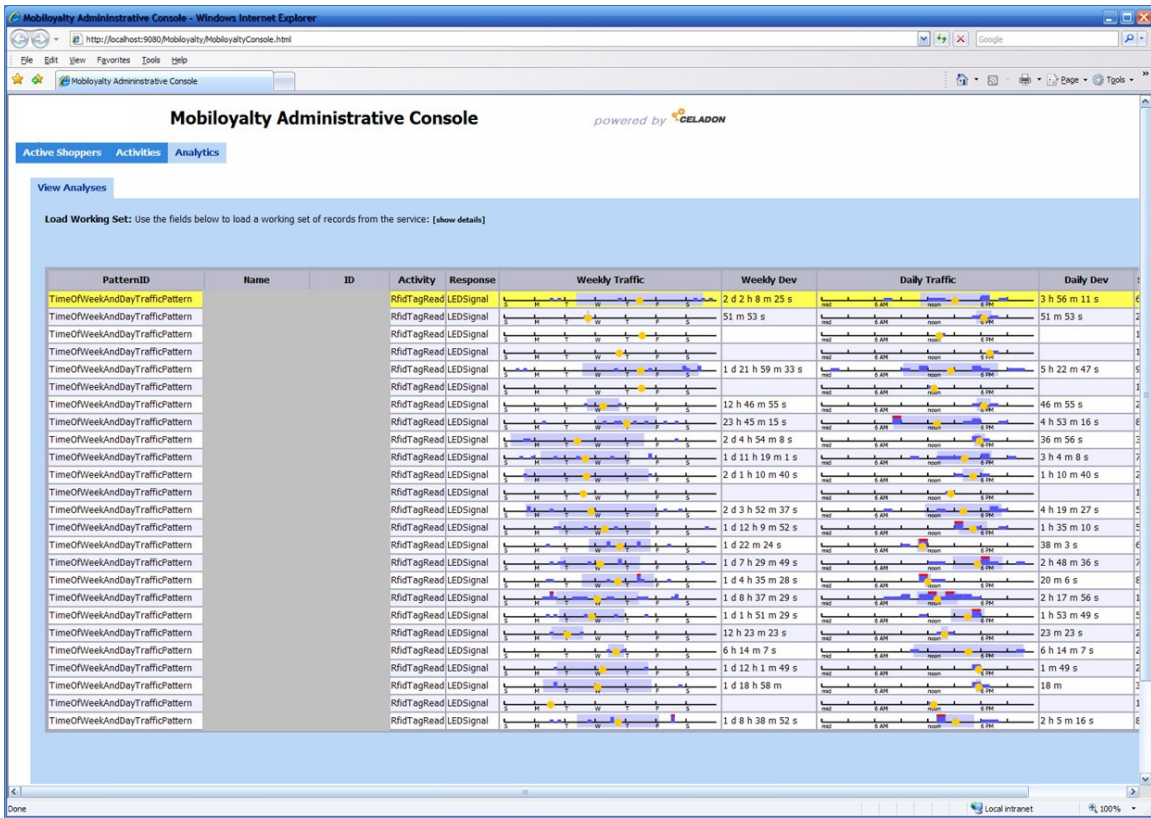


Figure 2. Celadon/Mobiloyalty traffic analyzer, multi-user view.

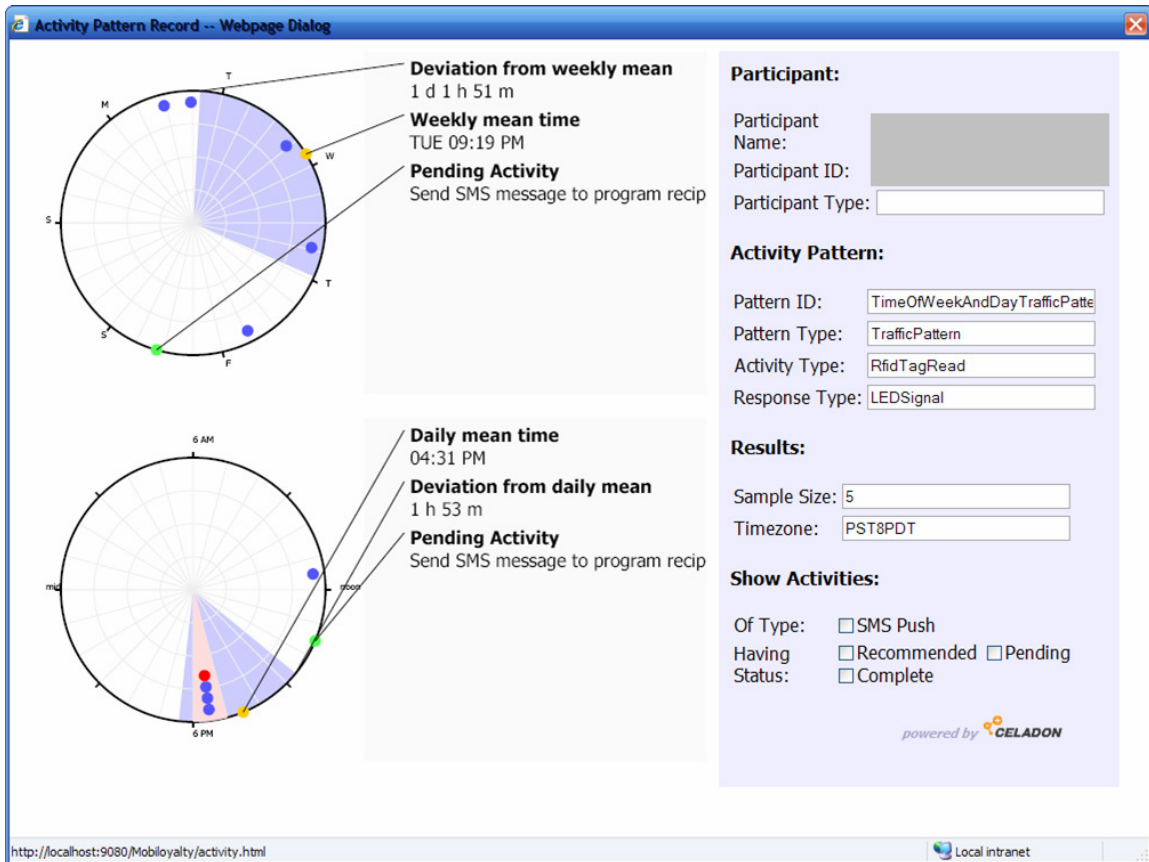


Figure 3. Mobiloyalty/Celadon campaign analyzer for a single user. The light red wedge indicates a high traffic rate for the 5-6 PM time slice.

3.5 Pluggable Architecture.

As noted above, Mobiloyalty builds upon a base technology, called Celadon, which provides a foundation for building mobile commerce applications and enabling them inside public spaces such as retail stores. The Celadon vision is seen in Figure 4. Celadon is a joint development of IBM Research and IBM's Ubiquitous Computing Lab (UCL) in Seoul, Korea. Celadon provides core software modules which allow retailers and other public space operators to assemble solutions such as Mobiloyalty. Standard components are plugged in and configured to construct an overall solution (for example, the swipe-oriented RFID recognition described above is plugged into the store environment, whereas in other environments other forms of localization could be plugged in).

The pluggable nature of the architecture also enables a distributed style of operation. In our pilot we used two servers, one in the store and one in the cloud. Both servers tracked customer context and state and made that information readable and writeable via REST, and also exposed events against the write operations. Both servers were loaded with a complete customer list.

The store server was used to operate the RFID reader point at which users tapped in. Due to its local customer listing, it was able to ascertain the validity of tap-ins, filtering redundancies (repeated taps), and quickly respond to the gesture with a sub-second response time. The response indicated the successful read of the NFC tag, and further indicated the availability of offers (or in cases that the user had already received their weekly offers, a "no new offers" feedback was given). The timing aspect is extremely important to the user experience, otherwise the user is unable to ascertain if

their actions were valid. The store server was used to partially drive the sequencing of the customer through the appropriate states, for example, executing the check offers response (Step C of Section 3.1). Because a copy of the customer's persistent state was maintained locally it was immediately possible to determine if the customer had been issued the applicable offers for the current offers cycle (one week). If the customer is sequenced into the "NewOffers" state, the store server also handled the initial message push to the device. This was performed via an attached GSM modem that could push an SMS message to devices via the local mobile telephone carrier.

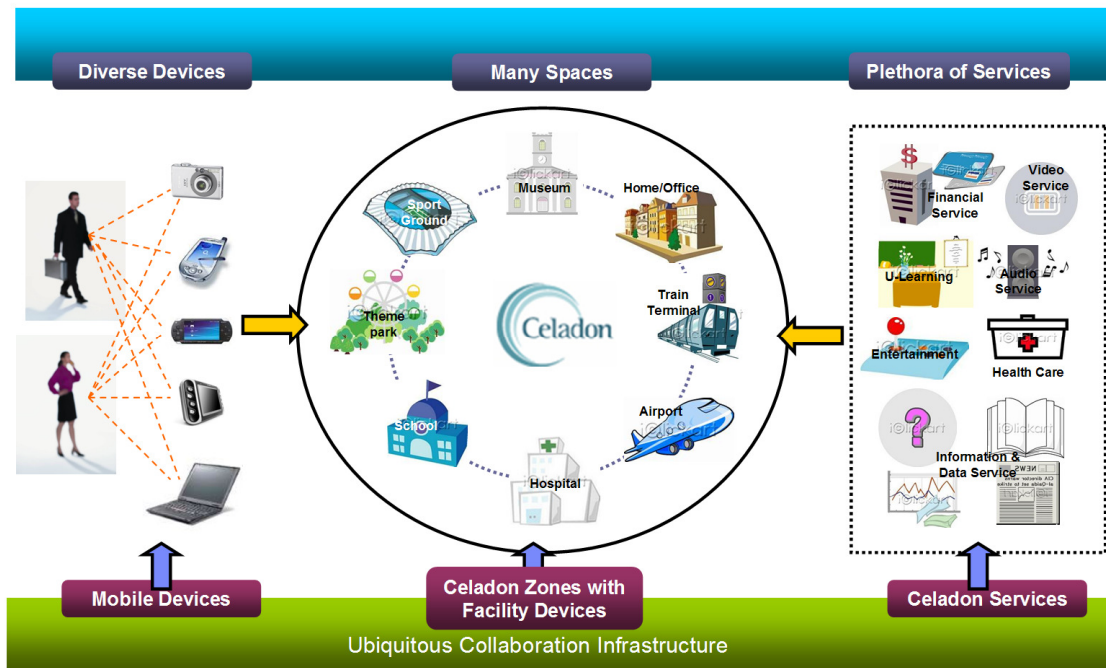


Figure 4. The Celadon technology vision for integrating mobile devices with mobile commerce services offered by merchants and other operators of public spaces.




Figure 5. The Celadon/Mobiloyalty management and monitoring center.

3.6 Cloud Operation.

The purpose of the cloud server was to provide user access to a more advanced web view of the customer's offers, perform analytics against the various activities at the stores, drive the periodic push campaign described in Section 2.3, and maintain a console infrastructure for the overall management of the program.

The cloud server maintained a periodically synchronous view of the local store server, although not at the subsecond level. The cloud server could provide synchronized results on a per-user basis via online consultation with the store server at any time, e.g., the user clicks through a link in an SMS message and starts their web application.

The cloud service also hosted the management consoles for each of the key solution elements. These consoles allow the solution to be monitored during its operations. Figure 5 illustrates the remote monitoring capabilities for the Mobiloyalty pilot, operated from our Hawthorne, NY location.

Console **a** is managing the mobile web service that provides HTML-based access for mobile devices equipped with web browsing capabilities. Console **b** is displaying a listing of active shoppers in the pilot store, their status, as well as arrival and expiration times. Console **c** displays individual activities within the system, including both “past” activities such as individual customer detection events (tap-ins), as well as “future” activities which have been scheduled such as SMS message pushes as part of the reminder campaign. Consoles **d** and **e** manage the analytics (see Figures 2 and 3). Console **f** is an advanced console, not used during the pilot phase, that displays alerts requiring operator attention, including customer activities throughout a store as well as status information for in-store systems. This console presents information graphically, correlating alerts to particular store locations, superimposed over an SVG diagram of the pilot store.

In general the system design struck a balance between edge (in-store) operations and cloud operations. We found that neither a pure edge-driven nor a pure cloud-driven design could fully address the realities of deploying this type of solution in a retail environment. Several key factors drove this design. First, there were potential latencies in the network traffic between the store and the cloud. In a pure cloud-driven design this could be a significant factor, especially when considering routing each tap-in to the cloud before making a response to the user. Thus, it made sense to perform all computations and immediate state transitions on the store server. Another factor was the need for a central point of management for the overall solution. This tilted the design away from a store-centric design. The centrality of the cloud is key especially in broadly cast solutions where users may travel from store to store and expect their state and context to follow them.

The cloud design also supports a multi-store or multi-tenant approach. This design provides the building blocks for such broader schemes involving multiple sites and schemes in which the cloud server hosts the user’s context and state and brokers that information to multiple players. However, additional research is needed into the development of robust and scalable value-added business models for managing, brokering and establishing the required policies and monetization schemes for such operations.

4: General Lessons Learned.

4.1 Simplicity of User Experience.

The project took a practical approach to channeling information to the user's device. SMS was used as the primary channel for extending coupon and promotion information to the device. Customers would receive an SMS message within a few seconds of swiping the tag enabled mobile device at the store entrance. The SMS message enumerated the featured offers available to the customer during their visit in summary form, due to the SMS message length restrictions of 160 characters. A URL link to a mobile device enabled page customized for the user is provided in the SMS message.. Upon selecting this link, mobile web-enabled customers link to a full web site which provides rich content,e.g. images and details, about the in-store offers.

The Mobiloyalty solution was an outwardly simple process for the customer, involving no more than a wave of the mobile device at the store entrance and reading a nearly simultaneous text message. In most cases the store experience was guided solely by the text message, few users clicked through to the richer web experience. Yet the user acceptance rates, even with a minimalist userinterface, were very high. This was not what our design and marketing teams anticipated. Drowning in a sea of rich data (a store full of 20,000 SKUs, devices capable of powerful GUI capabilities), how could a simplistic text message describing only a few items be effective? We concluded that the utter simplicity of the message tunnelled the attention space of the user, guiding them to the targeted products.

4.2 Time and Place Matter.

The Mobiloyalty solution capitalized on the user's time and place contextual information coupled with individualized knowledge of the retailer-customer relationship from the loyalty program. This resulted in an on-the-spot style of mobile marketing which paid off well. See below for acceptance rates.

Note that this approach could be viewed as a conservative mobile marketing strategy. Instead of attempting to engage the customer too far in advance of their store visit (time of demand) the system waits for a highly opportune time for engagement - store entry.

4.3 Normalization of Experience.

As the field of mobile marketing evolves, it is likely that marketers will recognize certain normalized patterns of opportunity in user shopping cycle. Normalization has already occurred implicitly in the traditional shopping cycle. There is one common point of interaction shared by almost all retail experiences - checkout and store exit.

One approach to mobile marketing is to focus attention on this part of the cycle as borne out by the proliferation of NFC projects focused on mobile payments. However, by careful attention to user patterns and expectations, additional interaction points such as store entry can be identified. We urge that the exploitation of these additional interaction points be normalized by the mobile marketing community so users know what to expect and do not feel spammed.

4.4 Engage Customers via Gesture-Reward.

The in-store experience establishes a gesture/reward style interaction with the customer. Via the voluntary tap-in at the reader point, the customer signals their desire to engage in a transaction with the retailer, collect information, or receive a reward. By using the tap-in methodology, it is possible for the system to establish fine-grain knowledge of where the customer is in the store, and what specific transaction the customer may be interested in. This approach is inexpensive, reliable, and simple, and is available to all device owners. The approach does not require elaborate physical configurations such as wireless signal triangulations. We also found via customer feedback that there is virtually no customer resistance to this methodology, and in particular, customers are made to feel that they are in charge of the process via their voluntary actions.

5: Findings and Conclusions.

5.1 Future Work: Extensions of the Concepts.

The scenario and components described in this paper can be extended in a variety of ways. Mobiloyalty uses a very simple approach to the management of the promotions on behalf of the user, there is only one list of promotions for each user, all offers are enabled and auto-consumed at the point of sale. This approach may be extended to a more generalized *digital wallet* [2]. As noted, there is a wide variety of alternate mechanisms for managing a digital wallet, including those in which the user has greater control of the arrangement and management of the items in the digital wallet.

The basic interchange with the user is based upon (1) monitoring user context, in particular location (2) awaiting trigger points based on changes in the context, user gestures, or schedules, and (3) customize an interaction based on the triggers and the user context. All aspects of this interchange can be generalized. For example, the user context is a combination of the user's location, their identity, their preferences, and their intentions. Much additional research is needed on each concept,

in particular how to establish preferences and intentions and match those to the retailer. A wide variety of triggering mechanisms are available, including different schemes for monitoring location (GPS-driven geo-fencing, detecting the network linkages of the device (cellular, WiFi, Bluetooth, etc.), as well as gesture driven forms such as RFID/NFC tagging, user scanning of visual patterns such as QR codes. Different trigger points and times could be used. For example, the tap-in at store entry could be generalized to other locations in the store (e.g. particular aisles, shelves, departments or service points) and associated with other functions. The trigger points could be generalized further to happen outside of the store, and be timed based upon marketing programs or aspects of the user's lifestyle. Additionally the nature of the interactions which are triggered can be generalized beyond that of pushing offers, for example, providing product information, advertisements or interactions which integrate the user with a business function such as package pickup. The opt-in reminders could be utilized as generalized opportunities for interaction beyond just reminding the user of promotions, by offering advertisements or other messages .

Additional research is needed into models in which all of the above aspects could be managed by third parties and brokered to retailers as needed. In our design, essentially all customer context and state information is provided directly to the retailers systems, and all interaction/response modules have global visibility across all customer data. In more refined models an ecosystem would be provided in which the retailer is a single player and customer information is filtered and brokered to its point of need.

5.2 Customer Acceptance Rates.

Another key goal of the Mobiloyalty project was to measure the inherent willingness of mobile device users to engage in mobile commerce functions using their device. Traffic and acceptance rates were tracked during the program and are summarized in Figure 5 below. This data is important in understanding the core willingness of customers to participate in programs such as Mobiloyalty, where their mobile devices may be used directly in the physical environments they are visiting. The program started with an email invitation sent to a candidate base (336) potential participants. The candidates, all employees of the retailer or IBM, were selected by the retailer based on the proximity of their home addresses to the pilot store. Approximately 20 percent (68) of the invitees responded to the invitation during a one week enrollment period. The customer received their electronic tag, attached it to their device, with a removable adhesive, and registered the mobile telephone number and loyalty identifier with the program. Of the enrollment base, approximately 66 percent (42) became active users in the program during the two month pilot program. Additionally, churn was tracked on a weekly basis. Approximately 41 percent of the active enrollees would visit the pilot store each week. Further, the program tracked the rate of conversion of the tap-in gesture to actual "transactions" -- in this case the redemption of the offers transmitted to the mobile device. The ratio of tap-ins to purchases of featured items was approximately 25 percent.

5.3 Anecdotal Observations and Issues.

There was generally a high level of curiosity and interest in the program, among the users who showed up for registration. There were no refusals of attaching the RFID tag to the phone and no serious concerns expressed regarding security ("Can someone secretly scan my tag and get information about me?"). Our team had anticipated that this might be a problem but found the actual experience to the contrary. In general mobile marketers may expect initial enthusiasm but should prepare for questions of this type *later on* the program as users think through privacy matters.

Our team maintained a fairly conservative set of expectations. Thus, the initial enthusiasm was suspect to us. This led our team to anticipate a trough of disillusionment, in which users simply forget about the program after enrolling and using the system once or twice. This expectation was the origin of the reminder campaign outlined in Section 2.3, and based upon our judgment,

implementation of the reminder campaign was key to maintaining a fairly consistent weekly traffic rate.

Finally, there were minor problems with the tags themselves. First, a very small number of tags were lost due to poor adhesive or being knocked off phones. Second, some mobile phones interfered with the radio frequencies interchange between the phone and readers. In these cases, we replaced the issued tag with a more rigorous "metal mount" tag which was imperceptibly thicker. The re-issued tags worked in all cases, and there were no cases in which no tag would work.

Our team managed the resolution of these problems by providing an email address to which customers reported problems. As the pilot was operated in a highly local environment it was easy to have the users simply stop by our store location and be refitted with a new tag. In general mobile marketing campaigns should be careful to provide some easy mechanism for problem resolution of this sort, and the problem resolution process should be scalable.



Figure 5. User acceptance and traffic rates measured during the Mobiloyalty in-store pilot, conducted summer 2008.

5.3 Conclusions

The above usage data should be viewed within the context of a small pilot, and the base numbers are relatively small. Thus the numbers at the end of the funnel should not be considered reliable. However, we found that weekly traffic was fairly consistent at about 40 percent. The redemption rate was based on a small base and should not be considered reliable – however when compared to redemption and usage rates of traditional paper coupons -- typically one or two percent would be considered strong – one may conclude that there is at least a great potential in the on-the-spot

offers paradigm. In summary then, assuming the usage rates hold when taken to a larger scale, the data are very strong and foreshadow tremendous demand by mobile device users for a wide array of mobile commerce functions.

Overall, these data bode well for customer acceptance of mobile device usage as an integral part of common retail transactions. It should also be noted that mobile devices open up substantial new opportunities that do not exist or have not been conceived, in which the mobile device can be used to construct new and powerful relationships between retailers and their customers.

References

1. McFaddin et al, Celadon: Delivering Business Services to Mobile Users in Public Spaces, IBM Research Report RC24381 (2007).
2. Cole, et. al, Toward a Mobile Digital Wallet. IBM Research Report RC24965 (2009).
3. Balan, R. K. and Ramasubbu, N. 2009. The Digital Wallet: Opportunities and Prototypes, IEEE Computer, Vol. 42, No. 4 (Apr. 2009), 100-102.
4. Bradford, T. and Hayashi F. 2007. Complex Landscapes: Mobile Payments in Japan, South Korea, and the United States, payments system research briefing, The Federal Reserve Bank of Kansas City, <http://www.kc.frb.org/PUBLICAT/PSR/Briefings/PSR-BriefingSept07.pdf>
5. Chaum, D. 1991. Numbers Can Be a Better Form of Cash than Paper, Computer Security and Industrial Cryptography 1991: 174-178.
6. Fling, B. 2009. Mobile Design and Development, 1st Edition. O'Reilly Media, Inc.
7. Hickson, I. et al. 2009. HTML5: A Vocabulary and associated APIs for HTML and XHTML, W3C WHATWG draft, <http://dev.w3.org/html5/spec/Overview.html>.
8. Irvine, M, Watches lose ground to cell phones, MSNBC Technology & Science, Feb. 16, 2007. <http://www.msnbc.msn.com/id/17189064/>
9. Landay, J., Joseph, A., and Reynolds F. 2009. Smarter Phones, IEEE Pervasive Computing, 8, 2 (Apr.-Jun. 2009, 12-13.
10. Mallat, N., Matti R., and Tuunainen, V.K. 2004. Mobile Banking Services, COMMUN ACM, 47, 5 (May 2004), 42-46.
11. McFaddin, S., et al. 2008. Modeling and Managing Mobile Commerce Spaces using RESTful Data Services. IEEE Intl. Conf. on Mobile Data Management, 2008.