



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Broadcast Scheduling for Mobile Advertising

Bert de Reyck, Zeger Degraeve,

To cite this article:

Bert de Reyck, Zeger Degraeve, (2003) Broadcast Scheduling for Mobile Advertising. Operations Research 51(4):509-517. <https://doi.org/10.1287/opre.51.4.509.16104>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

© 2003 INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

BROADCAST SCHEDULING FOR MOBILE ADVERTISING

BERT DE REYCK and ZEGER DEGRAEVE

*London Business School, Regent's Park, London NW1 4SA, United Kingdom
bdereyck@london.edu • zdegraeve@london.edu*

We describe a broadcast scheduling system developed for a precision marketing firm specialized in location-sensitive permission-based mobile advertising using SMS (Short Message Service) text messaging. Text messages containing advertisements were sent to registered customers when they were shopping in one of two shopping centers in the vicinity of London. The ads typically contained a limited-time promotional offer. The company's problem was deciding which ads to send out to which customers at what particular time, given a limited capacity of broadcast time slots, while maximizing customer response and revenues from retailers paying for each ad broadcast. We solved the problem using integer programming with an interface in Microsoft Excel. The system significantly reduced the time required to schedule the broadcasts, and resulted both in increased customer response and revenues.

Received February 2002; revision received January 2003; accepted January 2003.

Subject classifications: Marketing, advertising and media: mobile marketing. Production/scheduling, applications: advertisement broadcast scheduling. Integer programming, applications: IP solver interfaced with Microsoft Excel.

Area of review: OR Practice.

1. PRECISION MARKETING THROUGH MOBILE ADVERTISING

Marketing companies are concerned with targeting their campaigns to customers who might have or develop an interest in the marketed product or service. Advertising campaigns using traditional media typically generate low relative response rates. This requires targeting the masses, an expensive approach in terms of resources spent per actual customer. Precision marketing targets well-identified potential customers in advance, thereby increasing the response-to-advertisement ratio.

The increasing use of mobile phones has created a new opportunity for precision marketing. Barwise and Strong (2002) report a penetration of mobile phones in the United Kingdom of 70%, and up to 80% for young adults (18–24 years). Also, SMS (Short Message Service) text messaging is increasingly being used as a means of communication. According to various sources, more than 25 billion text messages are exchanged worldwide each month, with the United Kingdom being the second-largest European market, behind Germany. Barwise and Strong (2002) report that 68% of mobile-phone owners use text messaging, and up to about 95% of young adults. This means that approximately half of the U.K. population and three-quarters of all young adults can be reached via text messaging. Similar numbers are to be expected in other regions such as the United States, although the U.S. market has been slower to take off.

The high penetration of mobile phones and text messaging combined with the low cost of text messaging makes this an interesting medium for precision marketing. However, experiments have indicated that advertising via mobile phones works only if it is permission based. Barwise and Strong (2002) report that in the United Kingdom, 24%

of mobile-phone users would agree to receiving text-based advertising. Mobile advertising also allows sending advertisements to customers at a time when they are actually shopping, thereby increasing the chances that the ad will be acted upon. Moreover, the ads can be tailored to specific customers based on their profile, which contains information on their age, gender, lifestyle, and explicitly expressed preferences in terms of products or services.

Zagme, a company established in late 2000 by a graduate of London Business School's Sloan Program, is considered to be the pioneer in so-called location-sensitive wireless marketing, where customers receive advertisements depending on their location. At the time this project was carried out, *Zagme* was operating in two shopping centers in London, *Bluewater* and *Lakeside*, and was considering a quick expansion in the United Kingdom and Europe. *Zagme* acted on behalf of retailers wishing to attract customers to their stores with promotional offers, and had a contract with approximately 75% of the retail outlets in each of the shopping centers. In its first few months of operation, it had built a registered customer base of more than 80,000 people who had specified preferences for the different types of products and services that could be advertised. The products were classified in nine categories: Beauty, fashion, jewelry, gifts, sports, books, entertainment, restaurants, and miscellaneous. This information, complemented with age and gender information, was used to construct individual member profiles.

Customers announced their arrival at a shopping center by sending a text message to the company, logging them onto the system. From that time until they leave, again indicated by a text message, they received ads from retailers in the shopping center every hour on the hour. The control over whether or not ads are received was therefore

in the hands of the consumer, making it fully permission based. The rate of advertisements was restricted to one per hour to avoid saturation, which could lead to messages no longer being read but deleted upon receipt. A welcome ad was broadcast upon activation and a goodbye ad on deactivation. Consequently, on an average shopping trip of four hours, up to six messages would be received. Barwise and Strong (2002) indicate that 82% of the respondents of their survey on permission-based advertising mentioned that receiving three text messages per day was "about right," younger customers being more receptive to more frequent advertisements. However, because the messages in this case were broadcast when the customer was actually shopping, a higher number seemed acceptable or even desirable.

Although a few ads were brand building with a generic message, the vast majority were promotional offers of the direct response type, consisting of a discount or a free gift when a particular shop was visited within a specified time frame. The customer received the discount by showing the ad on the display of the mobile phone. Additionally, the customer received a monetary reward for each ad received, a few pence per ad, which could be used in the participating stores.

Early results indicated a great success, with some promotional offers causing a rush into the shop. Reebok experienced this using a loss-leader promotion consisting of a free pair of sneakers. More than 50 people ran into the store in the first four minutes after the offer was broadcast. A spokesperson for Reebok claimed that there were signs that this could become a key marketing tool. Customers typically spent between £10 (\$15) and £50 (\$75) each time they reacted to an offer, with response rates on average 10% and as high as 20%.

In addition to being an important new medium for marketing, mobile advertising also presents an emerging opportunity for operations research. Scheduling the advertisement broadcasts requires identifying which customers to target with which ads at what time, a complex task with a multitude of objectives and constraints. The problem was brought to our attention by the management of Zagme because they were scheduling the advertisement broadcasts manually, which was time consuming, tedious, error-prone, and preventing them from operating on a larger scale and expanding (inter)nationally. Advances in information technology enabled an integration of decision technology, marketing databases, and mobile-phone technology, allowing automated and optimized mobile advertising. In this paper, we describe such a new automated broadcast scheduling system for Zagme. In §2, we discuss the various aspects of creating mobile advertising schedules. In §3, we develop an integer programming formulation of the broadcast scheduling problem, and outline our solution approach in §4. Section 5 presents the implementation details, while in §6 we describe the results of using the system. In §7, we elaborate on the collaboration with the company and present a chronological overview of events. We discuss the system's limitations and directions for improvement in §8, while §9 contains our conclusions.

2. BROADCAST SCHEDULING

2.1. The Problem

Scheduling the broadcasts involved deciding which ads to send out to which active customers at what time. On the one hand, the company had a list of ads for which a retailer was willing to pay a preset amount if its ad(s) were broadcast. On the other hand, it had a list of active customers with different profiles. The broadcast schedule was constructed on a weekly basis; i.e., a schedule was generated for each upcoming week, two days in advance. Although schedules could be generated on a daily basis, the company considered weekly scheduling more practical. Also, it allowed taking into account which ads were broadcast on different weekdays, to ensure that customers who go shopping several days in one week receive different offers on different days. Each day was split into 12 time slots of one hour (10 am–11 am through 9 pm–10 pm), complemented by an extra activation and deactivation slot.

When the broadcast schedule was finalized, it was linked with the customer database, and a specialized system automatically broadcasted SMS text messages to the selected customers at the appropriate time. The message text was created by the company's marketing department, one for each advertisement and for each customer segment, and was highly customized depending on the targeted customer segment. The messages were designed in such a way as to entice a response, while fitting on the phone's display.

The company's profits were driven by revenues from the retailers. Hence, the broadcast schedule should be constructed in such a way as to maximize these revenues. Although short-term revenues can be maximized by considering the current set of available ads, potential future revenues depend on the size of the customer member base and the customer response rate. Hence, expected response rates would be a good proxy for expected future revenues. In the database marketing literature, a multitude of approaches are presented for response modeling, i.e., for determining the probability that a particular customer will make a purchase or not when he or she is targeted with a particular advertisement, in addition to the amount and the timing of the purchase. The approaches include logistic regression, linear and quadratic discriminant analysis models, neural networks, and Bayesian learning models (Baensens et al. 2002). Less research has been done when advertisements are customized for particular customers. A recent paper by Ansari and Mela (2003) describes a model for customizing communications and one-to-one marketing via e-mail using a Bayesian semiparametric approach for modeling customer response. Other papers discussing one-to-one marketing with response modeling include Rossi et al. (1996) and Shaffer and Zhang (1995).

Unfortunately, in this case the mobile-phone technology was not yet capable of recording whether a customer was effectively acting on the ad by purchasing the product or service advertised. Hence, actual response rates could not be measured directly, rendering the approaches discussed

above infeasible. Therefore, a proxy for expected response rates was used, based on the attractiveness of the offers broadcast and their appropriateness for a particular customer segment, assuming that attractive and well-targeted offers will result in an increased response. Because short-term profitability as well as long-term growth were both important objectives of the company, short-term revenues as well as offer attractiveness needed to be maximized. This required considering both the retailer perspective and the customer perspective, and constructing a broadcast schedule that balances their needs.

2.2. The Retailer Perspective

To maximize short-term revenues, priority was given to ads for which retailers were paying the highest fees. Although some retailers paid the standard fee, the company also provided its services to some retailers at a discounted cost, and sometimes even for free, in order to entice them to become clients. A star classification, 1* to 4*, was used to distinguish among clients. Typically, 4* clients paid the standard fee; 3* clients were paying a discounted fee; 2* clients were not paying at the time, but would most likely have started paying for the service soon, and finally a 1* was awarded to clients who probably were not going to become paying clients. These 1* clients were allowed free access to the system at nonpeak times in order to fill unused capacity and to provide customers with interesting offers at times where no alternatives from paying retailers were available. A higher broadcast priority was given to clients with a higher classification.

For each ad, the retailer could specify a preference regarding the timing of the broadcast. Because of limited capacity, the retailer was asked to provide three sets of time slots indicating a first, second, and third choice, as well as a list of other possible time slots in case the three choices were already taken, or to permit additional broadcasts if capacity allowed. A preference could be expressed for a contiguous time interval, e.g., Saturday between 2 and 8 pm, or for a noncontiguous set of time slots, e.g., every weekday at 11 am, or for a single time slot. If a preference was expressed for multiple time slots, this meant that the retailer was indifferent between the time slots in that set. Also, prebooking was allowed with advance payment and guaranteed broadcasts. Again, prebooking was possible for time intervals, noncontiguous time slots or single time slots. This enabled *Zagme* to guarantee broadcasts without losing too much flexibility. Obviously, prebooked slots were more expensive. In addition, for each ad a minimum and maximum number of broadcasts per week could be set. The minimum limit was used by the company to give retailers of relatively low importance a chance to experience the system as a preferred customer, and the maximum to prevent popular retailers being overrun by customers acting on promotional offers, and to increase the diversity of offers in the schedule.

The company used 12 different demographics or customer segments based on gender (M/F) and age (≤ 17 , 18–24,

25–34, 35–44, 45–54, ≥ 55). Most of the customers were in the age brackets ≤ 17 , 18–24 and 25–34, with a majority of women. Each retailer could request its ad to be broadcast to one or more appropriate customer segments.

2.3. The Customer Perspective

To maximize customer response and increase the size of the customer member base, the broadcast schedule should contain interesting customized offers that attract customer attention and entice an immediate response, while fostering loyalty and lock-in. Customer response will increase (1) with the *attractiveness of the ads*, i.e., whether or not they offer deep discounts, gifts, or offers on popular products; (2) with customized ads that *match the individual customer's preferences*; (3) with ads that are *received at an appropriate time*; and (4) with *the variety among ads* received. This will also result, through word of mouth, in a growing member base, potentially leading to higher prices charged to retailers.

Advertisement Attractiveness. Similar to the client classification, a star classification, 1* to 4*, was introduced to rate the attractiveness or quality of each ad from a customer perspective, with high ads' offers likely to generate a large response. Typically, 4* ads contained offers for free gifts or deep discounts on popular products, whereas 1* ads were generic brand-building messages. Priorities were assigned to ads based on their quality rating.

Matching Customer's Preferences. To match the personalized customer profiles, three ads of different types were scheduled in each time slot. If, given the customer's profile, the first ad was deemed inappropriate, it was blocked and the second one was broadcast instead. If the second ad was also inappropriate, a third one was broadcast. In effect, this meant that three different schedules needed to be constructed. The customer profiles were set up when a customer registered for the service via the Web, and could be modified at any time.

Broadcasting Ads at Appropriate Times. Retailers normally specified appropriate time slots for their ads. If no such preference had been expressed, the company's broadcast planners ensured that ads were broadcast at appropriate times.

Variety Among Ads. To achieve variety among the ads, the same ad was prevented from being broadcast more than once per day. In addition, two different ads of the same type were prohibited from being sent out in adjacent time slots in order to prevent repetition or broadcasting competing ads. Also, the same ad could not be broadcast in identical time periods on adjacent days because the company had observed that many customers shop on consecutive days at a similar time, especially on Saturdays and Sundays. Finally, diversity was ensured among the ads broadcast to different customer segments, enabling groups of people shopping together to receive different ads if they belonged

to a different segment. Finally, the ad in the first schedule was prevented from being reused in the second or third schedule around the same time by imposing a minimum three-hour time gap.

2.4. Balancing Retailer and Customer Preferences

Considering retailer- and customer-related objectives simultaneously resulted in a multiobjective decision problem. Multiple objectives are traditionally handled in one of two ways, either by creating a weighted sum of the different criteria to obtain a single objective, or by specifying a priority among the objectives, resulting in a multistage problem where each objective is optimized separately in the correct sequence. Our approach used a combination of both techniques.

Through discussions with management and the people responsible for scheduling, we were able to construct a priority list depicted in Table 1, determining which ads should be given priority depending on both the retailer and offer quality. This way, neither offer quality nor client quality had priority over the other, but 16 combinations of client and offer quality were determined, each with its own priority. As an example of an implicit trade-off, observe that 4* offers of a 3* client are given priority over a 3* offer of a 4* client, indicating that in this instance the quality of the offer is considered more important than the revenues from the retailer. However, the same is not true for 2* clients versus 1* clients, where a 3* offer of a 2* client gets priority over a 4* offer of a 1* client. It is important that these trade-offs and priorities are correctly determined, as they significantly influenced the resulting broadcast schedule. Naturally, the priority list could be amended at any time.

The time preference expressed by the retailers was handled similarly. Ads were assigned their preferred time slot as much as possible, with the first preference having priority over the second and the third. However, revenues and offer quality had a higher impact on priorities. For example, we would rather assign a second-preference time slot to a 4* client than a first-preference time slot to a 3* client if this would result in the 4* client being assigned a third-preference time slot. An example is provided in Table 2 below. Suppose that 10 ads are competing for the same time

Table 2. An example showing how priorities are determined based on client quality, offer quality, and time preference.

Client Quality	Offer Quality	Preference ¹	Priority ²
4*	1*	P	1
4*	4*	2	2
3*	4*	1	3
3*	4*	3	4
4*	3*	1	5
2*	3*	1	6
1*	3*	1	7
1*	3*	3	8
2*	2*	2	9
1*	1*	1	10

¹Low value means high preference. ("P" indicates a prebooked time slot.)

²Low value means high priority.

slot. Depending on client quality, offer quality, and preference expressed for that particular time slot, a priority value was determined.

3. PROBLEM FORMULATION

The system optimizes the broadcast schedule using integer programming with an interface in Microsoft Excel. We will first describe the formulation of the problem.

Parameters

O : Set of ads, index $o = 1, \dots, |O|$, $|O|$ is set at 100,

T : Set of ad types, index $t = 1, \dots, |T|$, $|T|$ is set at 9,

X_t : Set of ads of type $t = 1, \dots, |T|$,

p_{odscp} : Priority value for ad o broadcast on day $d \in \{1, 2, \dots, 7\}$ in time slot $s \in \{1, 2, \dots, 14\}$ to customer segment $c \in \{1, 2, \dots, 12\}$ in schedule $p \in \{1, 2, 3\}$. Each day of the week is numbered consecutively starting Monday, the time slots per day are numbered 1 through 14, with 1 and 14 denoting the activation and deactivation time slots, and the customer segments are numbered 1 through 12 denoting $F \leq 17$, $F18-24, \dots, F \geq 55$, $M \leq 17$, $M18-24, \dots, M \geq 55$.

The priority values are based on client quality and offer quality, using Table 1. Priorities are enforced in the objective function by multiplying the relevant decision variables with appropriate coefficients. The coefficients are set in such a manner that a choice with lower priority, if enforced in the schedule, will result in a lower objective value if it forces a choice with a higher priority out of the schedule. Also, the coefficients are increased or decreased in order to prioritize the client's time preference. Instead of determining appropriate objective function coefficients to obtain the desired effect, other approaches can be used, such as goal programming (Ignizio 1976), where different objectives are considered in sequence instead of simultaneously. For a review of different approaches for multicriteria decision problems, we refer to Keeney and Raiffa (1976) and Steuer (1986).

Table 1. An ad priority list determines the priority of ads based on client and offer quality.

Client Quality	Offer Quality	Priority ¹
4*	4*/3*/2*/1*	1/3/8/12
3*	4*/3*/2*/1*	2/5/9/13
2*	4*/3*/2*/1*	4/6/11/14
1*	4*/3*/2*/1*	7/10/15/16

¹Low value means high priority.

Decision variable

$x_{odscp} = 1$, if ad o is scheduled on day d , time slot s , to segment c in schedule p , 0 otherwise.

Objective: Maximize priority;

$$\text{Max} \sum_{o=1}^{|O|} \sum_{d=1}^7 \sum_{s=1}^{14} \sum_{c=1}^{12} \sum_{p=1}^3 p_{odscp} x_{odscp}. \quad (1)$$

Capacity constraints: Broadcast at most one ad in each time slot for each customer segment for the three schedules;

$$\sum_{o=1}^{|O|} x_{odscp} \leq 1 \quad d = 1..7; s = 1..14; c = 1..12; p = 1..3. \quad (2)$$

Intraday ad repetition constraints: An ad can only be broadcast at most once a day;

$$\sum_{s=1}^{14} x_{odscp} \leq 1 \quad o = 1..|O|; d = 1..7; c = 1..12; p = 1..3. \quad (3)$$

Interday ad repetition constraints: An ad cannot be broadcast in the same time slot on consecutive days (only for the first schedule);

$$x_{odsc1} + x_{o(d+1)sc1} \leq 1 \quad o = 1..|O|; d = 1..6; s = 1..14; c = 1..12. \quad (4)$$

Consecutive-type repetition constraints: Two identical-type ads cannot be broadcast in consecutive time slots (only for the first schedule);

$$\sum_{o \in X_t} x_{odsc1} + \sum_{o \in X_{t-1}} x_{od(s+1)c1} \leq 1 \quad d = 1..7; s = 1..13; c = 1..12; t = 1..|T|. \quad (5)$$

Interschedule-type repetition constraints: The three ads broadcast in a time slot for one customer segment should be of a different type;

$$\sum_{o \in X_t} \sum_{p=1}^3 x_{odscp} \leq 1 \quad d = 1..7; s = 1..14; c = 1..12; t = 1..|T|. \quad (6)$$

Interschedule ad repetition constraints: An ad cannot be scheduled within three hours of the broadcast of the same ad in any of the two other schedules;

$$\sum_{p=1}^3 \sum_{w=0}^2 x_{od(s+w)cp} \leq 1 \quad o = 1..|O|; d = 1..7; s = 1..12; c = 1..12. \quad (7)$$

Demand constraints: Each ad to be broadcast between a minimum and a maximum number of times (only for the first schedule);

$$l_{oc1} \leq \sum_{d=1}^7 \sum_{s=1}^{14} x_{odsc1} \leq u_{oc1} \quad o = 1..|O|; c = 1..12. \quad (8)$$

Secondary objective: Maximize intercustomer segment diversity (only for the first schedule);

$$\text{Min} \max_{0=1..|O|; d=1..7; s=1..14} \left\{ \sum_{c=1}^{12} x_{odsc1} \right\}. \quad (9)$$

The secondary objective minimizes the occurrences where an ad is broadcast at the same time to different customer segments. $\sum_{c=1}^{12} x_{odsc1}$ represents, for each ad, how many times it is broadcast to different customer segments at a particular time. One such value is obtained for each ad and time instant. The objective function (9) minimizes the maximum of $\sum_{c=1}^{12} x_{odsc1}$.

4. SOLUTION METHODOLOGY

4.1. Decomposition of the Three Schedules

To reduce the scheduling time, our system generates the three schedules sequentially instead of simultaneously, with the results of the previously generated schedule(s) acting as a constraint on the solution space of the new one. By relaxing Constraints (6) and (7) linking the different schedules, we obtain three separate problems resulting in a reduction in problem size. Obviously, this might result in suboptimal solutions as evaluated by the objective function (1). However, Equation (1) assumes that the three schedules are of equal importance, whereas the first schedule is relevant for many more customers compared to the second schedule, and even more compared to the third. Therefore, optimizing the three schedules in sequence is warranted because of the difference in importance in terms of customer impact and generated revenues.

Constraints (6) and (7) are enforced by reducing the available ads in the second and third schedule depending on the ads broadcast in the first schedule. Specifically, for the second schedule and a specific time slot, all ads of the same type of the ad broadcast in that time slot in the first schedule are eliminated. Also, the ads broadcast in the first schedule up to three hours earlier or later are removed from the set of available ads. Similarly, the available ads for the third schedule are constrained by both the first and second schedule.

4.2. Decomposition in Customer Segments

We generate the schedules for each customer segment separately. The only link among different customer segments is intercustomer segment diversity maximization in Expression (9). By eliminating this objective, a drastic reduction in problem size can be obtained. Intersegment diversity maximization can still be achieved heuristically by scheduling the customer segments sequentially and maximizing diversity with the customer segments already generated. We schedule the broadcasts for each customer segment separately, starting with the most heavily targeted customer segments. We compute for each potential ad a “diversity” value based on the ads scheduled in the already scheduled customer segments. More weight is placed on closely

5. IMPLEMENTATION

Observe that in the partial schedule in Figure 2, in each time slot three ads are scheduled as enforced by Constraint (2) in the problem formulation, except for the activation time slot, which is not fully utilized due to the fact that only a limited number of ads can be used for that

6. RESULTS

The goal of the development of the automated broadcast scheduling system was twofold. First, the company had a need for a system that would be able to develop schedules fast. The scheduling was done manually by a team of four people working for almost two days to develop a one-week schedule for the two shopping centers. Since the manual scheduling time is directly proportional to the number of

[illegible]

Figure 2. An illustrative example of part of a Monday broadcast schedule.**MONDAY**

	Female 17 or less	Female 18-24	Female 25-34	Female 35-44	Female 45-54
Activate	All Sports - SP	All Sports - SP Yo! Sushi 50% day - RE	All Sports - SP Yo! Sushi 50% day - RE	All Sports - SP	All Sports - SP
10.00	GT Recollections 1 - MI Pontis - RE Quicksilver 1 - SP	PizzaExpress 1 - RE Top Shop - FA The Bonsai House - MI	Dome Bar Café Coffee - RE Wallis - FA Quicksilver 6 - SP	Lush A - BE Dome Bar Café Coffee - RE Artworld - MI	Suits You - FA Artworld - MI Pontis - RE
11.00	Dome Bar Café Coffee - RE Top Shop - FA GT Recollections 2 - MI	Lush A - BE Dome Bar Café Coffee - RE Quicksilver 5 - SP	Pilot - FA Pontis - RE Waterstones Britney - BO	GT Recollections 1 - MI Giant Clothing - FA Pontis - RE	Lush A - BE Waterstones Travel - BO The Bonsai House - MI
12.00	Lush A - BE Dome Bar Café Meal - RE World of Football - SP	Yo! Sushi 20% - RE Watch It - JE Warehouse - FA	Lush A - BE Top Shop - FA World of Football - SP	Quicksilver 6 - SP Mikey - JE L'occitane - BE	Dome Bar Café Coffee - RE GT Recollections 2 - MI Club Golf - SP
13.00	Big Blue Rock - SP L'occitane - BE Warehouse - FA	Quicksilver 6 - SP Giant Clothing - FA Waterstones FPD - BO	Yo! Sushi 20% - RE Watch It - JE L'occitane - BE	Yo! Sushi 20% - RE Wallis - FA GT Recollections 2 - MI	Quicksilver 6 - SP Dome Bar Café Meal - RE Wallis - FA
14.00	Yo! Sushi 20% - RE Giant Clothing - FA All Sports - SP	Pontis - RE All Sports - SP L'occitane - BE	Quicksilver 6 - SP Dome Bar Café Coffee - RE Warehouse - FA	Mish Mash - FA Artworld - MI Dome Bar Café Meal - RE	Yo! Sushi 20% - RE L'occitane - BE Artworld - MI
15.00	GT Recollections 2 - MI Quicksilver 5 - SP Watch It - JE	Quicksilver 5 - SP GT Recollections 2 - MI Wallis - FA	Giant Clothing - FA Quicksilver 4 - SP The Bonsai House - MI	The Bonsai House - MI World of Football - SP Giant Clothing - FA	Warehouse - FA Quicksilver 1 - SP Mikey - JE

shopping malls, this prevented the company from operating on a larger scale. Rapid growth was deemed critical for the company to become profitable, emphasized by the fact that Zagme was already laying out a strategy for national expansion. Secondly, the system should increase the quality of the schedule in terms of customer satisfaction and response. This should eventually, through increased prices, lead to higher revenues from retailers.

6.1. Speed

The IP Formulation (1)–(9) results in 352,800 decision variables and 235,584 constraints. The two decomposition approaches discussed in §4 reduced the problem size to 9,800 variables and 18,036 constraints with, however, 36 different IPs to be solved. On a 2 GHz PC, this takes approximately 15 minutes, allowing for an interactive use of the system to generate different schedules. Also, multiple shopping centers can be scheduled in parallel.

6.2. Quality

Because no data on actual customer response could be captured, a proxy was used for expected customer response based on the quality of the broadcast schedule. The quality of the schedule was assessed by specialists in the company, and mainly depends on the quality of the offers broadcast and the variety in the schedule, as defined above. Overall, our system resulted in more attractive offers being broadcast, more ads matching customer profiles, more ads broadcast at the retailers' preferred time, and an increased variety among ads broadcast to different customer segments. Other observed advantages were a guaranteed prevention of intra-day, interday, and interschedule repetition. Although the scheduling team tried to take similar constraints into consideration when manually creating a schedule, several were

typically overlooked. As an example, ads were sometimes broadcast when they were not supposed to be broadcast, leading to a rush to a particular store which was not prepared for this and refused to act on the offer, resulting in disgruntled customers and retailers. Earlier, this had resulted in several retailers terminating their contract with the company.

To compare the manual schedules with the automatically generated ones, we ran the system in parallel with the manual scheduling procedure. The following improvements were observed for the test case in the week of April 2, 2001.

Observed Improvements

- In the manual schedule, 27% of the time slots (out of 1,176) were allocated to retailers who had specified that time slot as a preferred time slot, i.e., either as a first, second, or third preference. The automated system effectively doubled this to 55%. In the manual schedule, 18%, 5%, and 4% of the time slots were allocated to retailers' first, second, and third preference, respectively. The automated system increased this to 38%, 9%, and 8%.

- In the manual schedule, 121 time slots, i.e., more than 10%, were left unused due to time constraints or oversight, resulting in unwanted time slots being ignored rather than allocated to retailers as extra slots. This resulted in customers expecting promotional offers but not receiving any. Naturally, the automated system avoids unused broadcast capacity completely.

- In the manual schedule, 17 ads were broadcast to customer segments for which the ad was not appropriate. In total, 26 time slots were affected, i.e., approximately 2.5%. When confronted with this issue, the company's schedulers claimed that due to time constraints, checking these

restrictions manually was deemed too difficult and time consuming.

- In the manual schedule, 48 of the ads broadcast, i.e., more than 4.5%, were of the same type as the ad previously broadcast to that customer segment. Also, in the manual schedule, on 11 occasions the same ad was broadcast on consecutive days in identical time slots.

- In the manual schedule, diversity among the ads broadcast to different customer segments simultaneously was largely ignored due to the complexity of the scheduling task. Typically, ads were scheduled at the same time for all appropriate customer segments. Only in isolated cases was diversity taken into account.

- Manually, no backup schedules were constructed except for a very basic one, and that was essentially constructed by shifting the first schedule forward in time. As a result, customers would regularly not receive any ad when they had opted out of one or more product/service categories.

7. MODEL USE AND HISTORY

Between April and September 2001, *Zagme* used the automated broadcast scheduling system, while initially also developing schedules manually in order to be able to compare the results. Interestingly, by using the system the schedulers were also learning how to develop better schedules manually, for instance, by increasing the diversity in the generated schedules. Typically, the schedules generated by the system were used as a baseline, with deviations in specific circumstances to cope with new criteria and idiosyncratic constraints. The Excel interface made it very easy to change the generated schedule manually. In the months that followed, we were in close contact with the head of the scheduling team in order to make small adjustments to the system if required.

In September 2001, *Zagme* struck a deal with Channel 5, the leading commercial television channel in the United Kingdom, to set up a joint venture featuring interactive TV ads with the possibility of text message responses, allowing advertisers to form personal relationships with viewers who respond, offering discount vouchers or product information via their mobile phones. This would be the first step towards national expansion.

Unfortunately, in the wake of the events of September 11, 2001 and the economic downturn, *Zagme* started to experience financial distress, and started looking for a merger or a takeover (as a target). As a consequence, the operations of the company were scaled back to a very low level to minimize expenditures. This resulted in a drastic reduction of the client base, and several key people, including some in the scheduling team, started to leave the company. This finally resulted in the bankruptcy of *Zagme* on October 24, 2001. Since then, other companies such as *The Mobile Channel*, described in Barwise and Strong (2002), have been set up using a similar business plan.

8. DIRECTIONS FOR IMPROVEMENT

Although the system provided a major improvement both in terms of scheduling time and quality, further improvements can be made. A faster system could enable real interactive scheduling, but this would require the response time to be in the order of seconds rather than minutes. We have experimented with different approaches, including a dedicated multilevel branch-and-bound algorithm. However, preliminary results indicated that the branch-and-bound algorithm could not significantly outperform the IP solver, at least not to reduce the time needed to the order of seconds. The greater flexibility of the IP approach, where changes to the IP formulation can be done quite easily, therefore, makes this approach more suitable, especially because requirements, constraints, and preferences for customers as well as retailers were not yet fully known and are constantly changing. For instance, if the company would decide that it would be advantageous to broadcast several offers simultaneously, the model can easily be modified by using the offers in the second and third schedule, if appropriate. If necessary, additional schedules can be generated. Because of the decomposition approach used to reduce the complexity of the model, the time required to generate additional schedules will only increase linearly with the number of schedules.

The broadcast schedule was established on a weekly basis and was fixed. A natural extension would be to allow real-time dynamic scheduling, where broadcast decisions are made dynamically depending on which customers are active and their actual profiles. Additionally, the system could be extended by asking the customers to specify at log-in how many ads they would like to receive and in what product types they are currently interested. This concept is in line with the Sense and Respond methodology described in Bradley and Nolan (1998), where the customers' needs are sensed electronically in real time, transmitted using the mobile-phone technology, and linked to the scheduling system, resulting in a dynamically optimized broadcast schedule designed to respond to those needs in an optimal manner. This would in effect result in interactive marketing, with two-way communication between customers and retailers.

One could also consider taking advantage of new upcoming technology in mobile telephony that allows tracking of the exact position of a customer to within a few meters, thereby creating opportunities to broadcast ads depending on the exact location of the customer. Naturally, this will work only in a permission-based approach. This would allow the system to be used outside the realm of shopping centers by sending an ad to a customer if he or she approaches a specific store. This transforms advertising from a push-based system to a pull system, where individual customers trigger ads to be broadcast to them depending on their location. This could lead to interesting new areas of direct marketing by enabling potential customers to be "lured away" from competitors, for instance, by broadcasting a promotional offer to a potential customer entering

a competing store. The ramifications of such an approach have not yet been fully studied, however. A recent book by Newell and Lemon (2001) discusses the opportunities and challenges created by these technological advances and proposes several new marketing strategies in this context.

It is clear that the emergence of mobile advertising creates a number of challenges, but also opportunities, to the operations research community. The complexity of precision marketing requires decision support systems that allow taking full advantage of the potential benefits created by targeting individual customers based on detailed customer profiles. Recent advances in information technology, in terms of linking decision support systems with marketing databases and mobile technology, now allow the development of sophisticated decision tools, allowing the full potential of emerging marketing opportunities to be realized.

9. CONCLUSIONS

Permission-based mobile marketing, or advertising via text, voice, or picture messaging, will undoubtedly become more important in the future as a means of direct marketing. Instead of mass advertising a product or service, mobile marketing allows targeting well-identified potential customers based on their current location, thereby increasing the response-to-advertisement ratio. Given the high penetration of mobile phones and the relatively low cost of sending text messages, this is an interesting medium for direct marketing. Also, the ads can be tailored to specific customers using customer profiles and can be broadcast at an appropriate time—for example, when the potential customer is actually shopping.

However, as well as creating enormous opportunities, mobile marketing also creates interesting new challenges. One challenge is to decide which ads to broadcast to which customers at what time, given all the information at hand. Doing this properly can ensure the long-term viability of a mobile marketing business. In this paper, we presented a new broadcast scheduling methodology for mobile marketing developed for *Zagme*, a company established late 2000 in the United Kingdom.

The problem, formulated as an integer program, is solved by a dedicated IP-solver with a user interface in Microsoft Excel. The model balances the needs and preferences of both customers and retailers, resulting in a multiobjective setting, solved by determining appropriate priorities. The system maximizes customer satisfaction primarily by broadcasting interesting offers such as deep discounts or

offers on popular products, customized ads that match the individual customer's preferences at an appropriate time, and with maximum variety among ads. Further, the system also considers retailer preferences by tailoring the ads to specific customer segments, to individual customer profiles, and by allowing the specification of preferred broadcast time slots. Moreover, the developed system greatly reduces the time required to schedule the ad broadcasts, allowing for an increase in the scale of operations.

ACKNOWLEDGMENTS

The authors thank Bill Green, CEO of *Zagme*, and Russell Buckley, marketing manager, for bringing this problem to their attention. They are also grateful to Helen Keegan and her broadcast scheduling team for evaluating and testing the newly developed system. The authors also thank the editor, associate editor, and two anonymous referees for their suggestions for improving the paper.

REFERENCES

- Ansari, A., C. F. Mela. 2003. E-Customization. *J. Marketing Res.* **40**(2) 131–146.
- Baesens, B., S. Viaene, D. Van Den Poel, J. Vanthienen, G. Dedene. 2002. Bayesian neural network learning for repeat purchase modelling in direct marketing. *Eur. J. Oper. Res.* **138**(1) 191–211.
- Barwise, P., C. Strong. 2002. Permission-based mobile advertising. *J. Interactive Marketing* **16**(1) 14–24.
- Bradley, S. P., R. L. Nolan, eds. 1998. *Sense and Respond: Capturing Value in the Network Era*. Harvard Business School Press, Boston, MA.
- Ignizio, J. P. 1976. *Goal Programming and Extensions*. Lexington Books, Lexington, MA.
- Keeney, R., H. Raiffa. 1976. *Decision Making with Multiple Objectives*. Wiley, New York.
- Newell, F. B., K. N. Lemon. 2001. *Wireless Rules: New Marketing Strategies for Customer Relationship Management Anytime Anywhere*. McGraw-Hill, New York.
- Rossi, P. E., R. E. McCulloch, G. M. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* **15**(4) 321–340.
- Schrage, L. 2000. *Optimization Modeling with Lingo*. Lindo Systems Inc., Chicago, IL.
- Shaffer, G., Z. J. Zhang. 1995. Competitive coupon targeting. *Marketing Sci.* **14**(4) 395–416.
- Steuer, R. E. 1986. *Multiple Criteria Optimization: Theory, Computation and Application*. Wiley, New York.