

# Extended Abstract: New Media Planning Models for New Media

John Turner

The Paul Merage School of Business, University of California at Irvine, Irvine, CA 92697-3125  
john.turner@uci.edu

---

Advertisers are always looking for new and exciting ways to catch peoples' attention. For example, when Massive, Inc., debuted its dynamic in-game advertising network in 2002, advertisers latched onto the idea that branded imagery could be dynamically inserted into video games. Although video games have incorporated static ads for decades, dynamic in-game advertising is new: the technology allows in-game billboards to display different ads to different players based on their demographic, time of day, day of week, and other criteria. Imagine playing a hockey game and seeing the logo for your local pizzeria on the arena's side boards, and this week's blockbuster movie advertised on the Jumbotron above. Not only do such ads add realism to video games, but they provide marketers with a way to reach young people that are watching less television.

In addition to dynamic in-game advertising, the entertainment industry has been working on embedding targeted ads – those shown only to audience segments requested by advertisers – in a wide spectrum of new media. For example, Visible World (WARC Newswire 2009) is piloting targeted TV technology that breaks the broadcast paradigm by allowing different ads to be served to different households that are concurrently watching the same TV show. Meanwhile, Zunavision (Ivanhoe Newswire 2009) is developing technology to dynamically paint logos on top of live video, so ads appear seamlessly within TV shows. And finally, Microsoft and Intel are developing an electronic billboard (WARC Newswire

2010) which will use a video camera and artificial intelligence to infer the age and gender of viewers so relevant ads can be displayed and exposures tallied.

Each day, billions of targeted ads are channeled to viewers via ad networks – intermediaries which package and sell ad space from multiple publishers’ websites, video games, or other media vehicles. To match ads to appropriate audience segments, meet advertisers’ goals, and ensure opportunities to serve advertising are not wasted, ad networks solve complex planning, scheduling, and pricing problems.

We begin in Section 1 by solving the scheduling problem faced by Massive Inc., a wholly-owned subsidiary of Microsoft and one of the first dynamic in-game ad networks. At the heart of Massive’s scheduling system is an **ad server**. Each time a player enters a new level of a game, his console (Xbox or PC) connects to this ad server via the Internet and requests new ad graphics for billboards, stadium walls, and other locations where ads may be placed in the level. The ad server decides which ads to serve to this player, and functions much like a web server delivering banner ads to web pages. But, unlike on the web, where selected ads are almost certainly seen, only a fraction of selected in-game ads will be seen by the viewer. For in-game advertising, billable ad time is recorded only when, as part of normal game play, the player navigates through the level and passes locations where ads are displayed. For this reason, and also because of additional constraints (such as **saturation**, **competition**, and **context**, discussed below), scheduling in-game ads is significantly more complicated than scheduling banner ads on web pages.

The operational problem in-game ad networks like Massive face is how to schedule and serve ads to players over time in a way that makes the best use of their inventory of ad space. Campaigns purchased by ad agencies specify a target number of **impressions** (ads seen by gamers), a rough schedule for serving these impressions over time, and also a desired mix (e.g., 60% in sports games, 40% in the rest of the games). A campaign’s delivery may also be restricted to certain geographic areas and/or times of the day. In addition, the network provider must also manage 1) **saturation**: it is undesirable for a single player to simultaneously see many copies of the same campaign, 2) **competition**: campaigns of two competing brands – e.g., Coke and Pepsi – should not be served to the same player, and 3)

**context:** ads should not seem out of place within the game – e.g., Coke ads belong on virtual soda machines, Tide ads do not.

The size, scope, and complexity of Massive’s problem are such that even if there were no system uncertainty, optimization of their ad server would be difficult. But of course, uncertainty is present – there are three primary sources: 1) the acquisition of new games, 2) the sale of new campaigns, and 3) error in inventory forecasts of ad space. This last factor, uncertainty in the amount of ad space, arises because the number of players, the types (demographics) of players, and the ad space that the players actually see during game play are not known when the scheduling problem needs to be solved. Thus, campaigns sometimes fall short of their impression goals or deviate from the desired pattern of delivery; in that case, the network provider offers the advertiser a **make-good**: the campaign is extended or the advertiser is offered a refund or credit for future use.

We present the first scheduling model and algorithm for dynamic in-game advertising. Our model has two components: 1) a linear program called the **Weekly Plan** is solved periodically to establish target service rates, and 2) a packing heuristic called the **Real-time Algorithm** is run whenever a player enters a new level, to serve impressions in accordance with these service rates while taking other side constraints into account. Benchmarking our model against Massive’s legacy algorithm using 26 weeks of historical data, we observe 1) an 80-87% reduction in make-good costs (depending on forecast accuracy), and 2) a shift in the age distribution of served ad space, leaving more premium inventory open for future sales. Moreover, when Massive made changes to its systems to exploit the key components of our plan-then-execute model, the number of unique individuals reached by each ad campaign increased by on average 26% while campaign delivery, as measured by the standard deviation of hourly impressions served, became 33% smoother.

Next, in Section 2 we build on our real-world experience and formulate a single-period ad planning problem which focuses on the core structure of how ad networks should plan advertising in a broad class of new media, which we call *Guaranteed Targeted Display Advertising* (GTDA). Included under the umbrella of GTDA are dynamic in-game advertising, webpages that display banner ads and video ads, social media platforms like Facebook that match ads to users’ profiles, as well as digital TV, dynamic

TV product placement, and intelligent electronic billboards. In general, we define GTDA as the class of advertising that satisfies the following four properties:

- CPM Sales Model: Advertisers pay for a number of “eyeballs,” called *impressions*. Each impression corresponds to an individual that sees an ad at a particular point in time, e.g., by viewing a banner ad on a webpage. Prices are quoted in cost-per-thousand (CPM); e.g., a \$30 CPM means \$30 buys 1000 impressions.
- Measurable Progress: The exact number of impressions served to date is known.
- Targeting Control: Ads shown to a specific individual can be chosen based on that individual’s characteristics (demographic, geographic, and/or behavioral).
- Guaranteed Delivery: The ad network promises to serve each advertiser an agreed-upon number of impressions over a fixed time period. In this sense, delivery is “guaranteed,” since ad networks do whatever they can to avoid under-delivery. Due to considerable uncertainty in the audience sizes of the various audience segments, ad networks must choose carefully when deciding which individuals get served which ads.

We solve the ad network’s single-period GTDA planning problem that allocates impressions generated by multiple audience segments to multiple ad campaigns. Since impressions are guaranteed and negotiated in advance, the ad network’s principal concern is to provide a high quality of service to its clients, the advertisers. Therefore, the ad network prefers plans that yield both *high reach* and *low variance*. High reach means a large number of *unique* individuals see each ad campaign. Low variance means the number of impressions actually served by executing the plan is highly predictable. By explicitly modeling audience uncertainty, forecast errors, and the ad server’s execution of the plan, we derive sufficient conditions for when the solution to the GTDA plan minimizes variance and maximizes expected reach.

From a computational standpoint, GTDA ad networks must solve planning problems with millions or even billions of audience segments. This arises because the subsets of viewers that different advertisers target can intersect in complex ways. Therefore, even when each advertiser targets only a handful of customer segments, the ad network must plan at a much finer resolution to account for the interactions

among the full set of advertisers that have purchased intersecting blocks of audience. For example, if advertisers can choose to target 2 genders, 10 age categories, and 5 income levels in 500 geographic regions over 7 days and 6 dayparts in any conceivable way, the planning problem has  $2 \times 10 \times 5 \times 500 \times 7 \times 6 = 2.1$  million viewer types. Yet over a short time window, impressions only come from a sparse subset of these viewer types. To address this issue, we develop two efficient algorithms for intelligently aggregating the audience space: the first assumes management specifies a fixed clustering of audience segments to use in an aggregate problem, and appropriately scales the aggregate solution to produce a feasible solution in the disaggregate space; the second intelligently refines the audience space into successively smaller clusters, converging to an optimal solution for the original disaggregate problem without ever explicitly solving the disaggregate problem. Often, near-optimal schedules can be produced despite significant aggregation.

To summarize, we contribute to the ad planning literature by defining GTDA as a specific class of targeted advertising which cross-cuts various types of media and is characterized by four structural properties that allow ads to be planned in a common manner. We show that for GTDA, minimizing a quadratic objective is a good surrogate for maximizing reach and minimizing variance, and derive sufficient conditions for when reach and variance are exactly optimized. Furthermore, we introduce two algorithms to intelligently aggregate the potentially large audience space; these algorithms exploit the spreading which our quadratic objective induces.

Finally, in Section 3, we study the ad network's problem of pricing targeted advertising. Our motivation is to generate managerial insights, and so we study several models where an ad network sells ad space to two heterogeneous advertisers. We characterize the optimal solutions under perfect information for two cases: one where advertisers buy as much advertising as they can get for a reasonable price, and another where advertisers always purchase according to their given spending targets. In addition, we study the effect of advertisers' unknown willingness-to-pay under several advertising objectives, such as advertising as much as possible to build brand image and advertising a set amount to clear current-period product inventory.

## References

Ivanhoe Newswire. 2009. The Future of TV: Hidden Commercials? Ivanhoe Newswire, July 2009. Available at <http://www.ivanhoe.com/science/story/2009/07/599a.html>.

WARC Newswire. 2009. Google offers targeted TV ads. WARC Newswire, 29 July 2009. Available at <http://www.warc.com/News/PrintNewsItem.asp?NID=25464>.

WARC Newswire. 2010. Microsoft, Intel develop interactive digital signs. WARC Newswire, 14 January 2010. Available at <http://www.warc.com/News/PrintNewsItem.asp?NID=26183>.