

Symphony: Enabling Search-Driven Applications

John C. Shafer

Search Labs, Microsoft Research
1065 La Avenida St.
Mountain View, CA 94043

jshafer@microsoft.com

Rakesh Agrawal

Search Labs, Microsoft Research
1065 La Avenida St.
Mountain View, CA 94043

rakesha@microsoft.com

Hady W. Lauw

Search Labs, Microsoft Research
1065 La Avenida St.
Mountain View, CA 94043

hadylauw@microsoft.com

ABSTRACT

We present the design of Symphony, a platform that enables non-developers to build and deploy a new class of search-driven applications that combine their data and domain expertise with content from search engines and other web services. The Symphony prototype has been built on top of Microsoft's Live Search infrastructure. While Symphony naturally makes use of the customization capabilities exposed by Live Search, its distinguishing feature is the capability it provides to the application creator to combine their proprietary data and domain expertise with content obtained from Live Search. They can also integrate specialized data obtained from web services to enhance the richness of their applications. Finally, Symphony is targeted at non-developers and provides cloud services for the creation and deployment of applications.

1. INTRODUCTION

While people world-wide now turn to a general web search engine to satisfy their general information needs, they at the same time seek specialized sites for information on specific topics. Examples include Expedia for travel, Amazon for books, eBay for rare artifacts, WebMD for health, IMDB for movies, Ticketmaster for events, etc. Google Custom Search (google.com/coop/cse), Yahoo! BOSS (developer.yahoo.com/search/boss), and Live Macros (live.com/macros) represent efforts by general search engines to enable people to build custom search engines that alter the default behavior of the underlying general-purpose search engine. Such a custom search engine may, for example, restrict the search to some domains, automatically add terms to an input query, or reorder search results to give preference to some URLs.

A natural progression in this line of evolution is for a search engine to take up the role of a world-wide resource for unstructured data to be used in building and deploying search-driven applications. For instance, a video store owner may be interested in an application that allows people to browse his inventory of movies, but augments his content by integrating it with focused search results for supplemental content such as the latest reviews and trailers obtained on the fly from a search engine. A wine connoisseur may create and embed in her web site a specialized search vertical that combines her knowledge of wines with targeted web-search results, and may be able to

monetize her efforts using an advertising or referral service.

We present the design of a system, christened Symphony, to assist in the development of such search-driven applications. While Symphony naturally makes use of the customization capabilities offered by the search engines (our prototype uses Live Search), its distinguishing feature is the capability it provides to the creators of the search applications to be able to combine their propriety data and domain expertise with the content obtained from the search engine (and other data services). In this way, Symphony users can focus on developing and curating quality content, making the best use of their specialized interests and knowledge, while still being able to fall back on a search engine for general content. The integration of specialized data and Web resources makes richer results possible.

Another important feature of Symphony is its target audience. Rather than professional application developers, Symphony is designed to be usable by domain experts and enthusiasts, who may not have the requisite development skills to build a search application from scratch. Symphony carries out the heavy lifting, and allows its users to concentrate on functionality and design. The key capabilities provided by Symphony include provision for storage and indexing, integration of proprietary data with Web resources, easy-to-use design interface, support for monetization, and deployment of search applications to users' Web sites or social network platforms (e.g., Facebook).

2. SYSTEM DESIGN

We present here the key aspects of Symphony's design. We first describe Symphony's functionality from the perspective of an application designer and then describe a few system internals.

2.1 Creating a Search-Driven Application

Assume for present purposes that the application designer already possesses the proprietary content and domain expertise to be incorporated into the application. Content could be inventory data for a retailer, a list of favorite movies or wines for an enthusiast, etc.

Symphony offers the following capabilities to the designer.

Proprietary Data: Symphony provides private and secure space to store and index proprietary data belonging to the application designer. It supports a variety of upload methods (e.g., HTTP/FTP file upload, RSS feeds, or URL crawling), as well as a variety of structured data formats (e.g., delimited files, Excel files, and XML). Symphony also supports dynamic data accessed through SOAP and REST-based web services. This facilitates real-time data freshness, allows users to keep data considered too sensitive "in-house" and allows integration of 3rd-party services.

Built-in Services: Symphony has built-in support for accessing internet content in the form of web/image/video/news results from a search engine. Most services support additional configuration,

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Database Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permissions from the publisher, ACM.

VLDB '09, August 24-28, 2009, Lyon, France.

Copyright 2009 VLDB Endowment, ACM 000-0-00000-000-0/00/00.

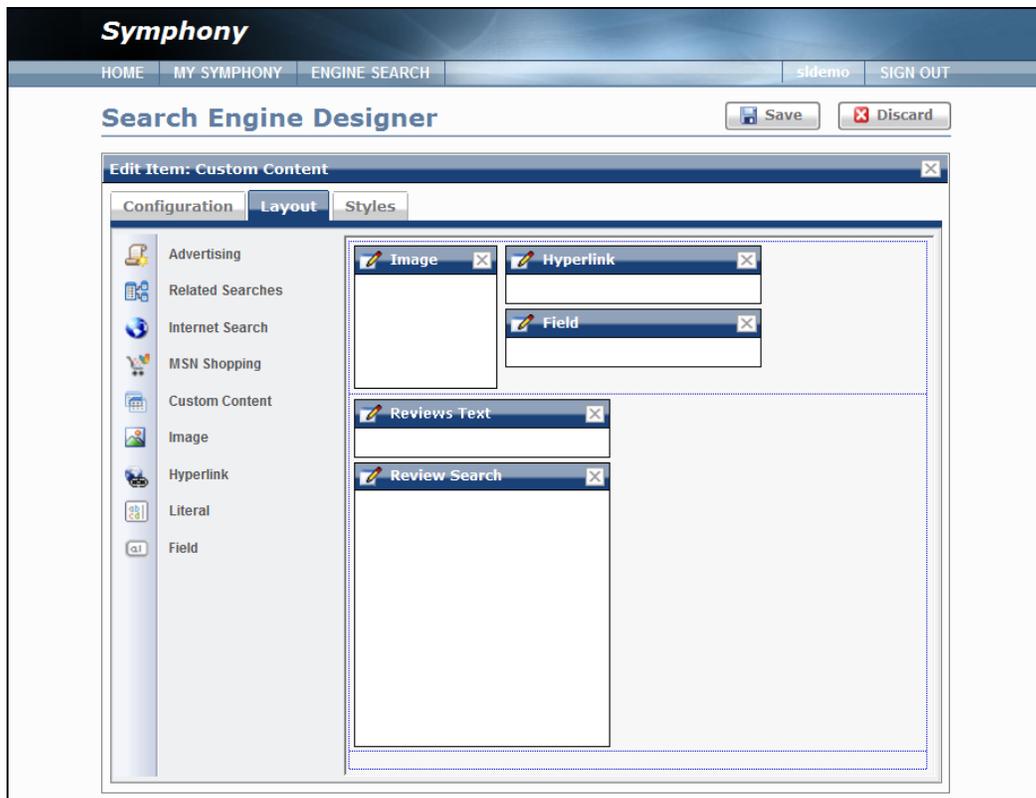


Figure 1: Design Interface

such as site restriction. A Site Suggest [2] feature is provided that can suggest additional related sites to include based on the set already specified. We also integrate with advertising services such as adCenter (adcenter.microsoft.com), allowing ads to be displayed and configured just like any other content source.

Data Integration: The various proprietary, 3rd-party and built-in data sources can be integrated flexibly, for instance by using one data source as a primary content (which will be queried based on the application user's query) and one or more data sources as supplemental content (which will be queried based on selected fields from the primary content).

Design Interface: Symphony provides a WYSIWYG-style design interface – no coding is required. Figure 1 shows a screenshot of the design interface implemented in our current prototype. The left bar shows various data sources that application designers can drag-n-drop onto an application, specifying how each should be searched and how many results to be shown. Multiple data sources can be added to the layout and arranged as desired. This drag-n-drop process is also used to configure how individual results should be laid out. Application designers can create HTML elements such as text, images and hyperlinks using fields from the data source. The right panel in Figure 1 illustrates one such layout. In this case, a search result features a hyperlink, an image, and a descriptive field. In addition, a search result may trigger another Internet search over supplemental content. Supplemental content can be added by simply dragging additional data sources onto the current result layout. The designer selects which fields from the first data source to use when querying that secondary data.

Presentation: In addition to general layout organization, further customization of the application's look-and-feel is supported via templates, wizard-style assistance from Symphony, or through style properties on individual elements (e.g., color, font-size). For more web-savvy users, greater control is possible via style-sheets.

Distribution: Application designers can embed the application into their own web sites by copy-and-pasting auto-generated snippets of JavaScript and HTML onto a web page. Symphony also supports publishing the application to social networking sites such as Facebook.

Hosting: Regardless of how an application is distributed, its execution and the resources involved are always shouldered by Symphony. This facilitates seamless integration while minimizing any resource requirements on the part of the application designer.

Monetization: Symphony has built-in support for the application designer to be able to record customer interactions with the application and obtain various summaries. When a link is clicked in a Symphony-hosted application, it can be logged by the system. If the click is on an advertisement from an integrated ad service, the application designers will automatically be credited by that service for any ad-click revenue. In another scenario, a summary of an application's click traffic can be downloaded by the application designer to serve as the basis for charging or auditing referral compensation.

2.1.1 An Application Example

A video game store owner, Ann, wants to create a custom search experience around her inventory of video games. After registering her proprietary inventory data with Symphony, Ann drags the inventory data onto a new application layout as a primary content,

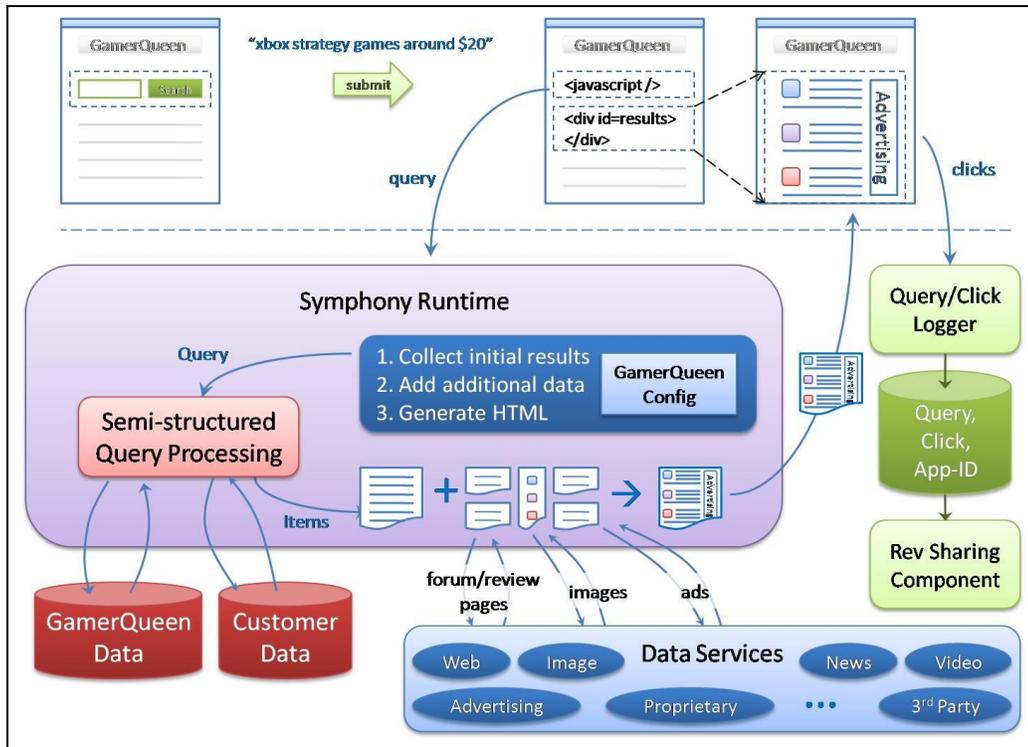


Figure 2: Query Execution in Symphony

and configures the application to search by title, producer, and description. She then configures the result layout to show the title hyperlinked to a detail page, an image, and a description. Ann may then wish to include game reviews as supplemental content by dragging web-search content onto the result layout and restricting it to sites such as gamespot.com, ign.com and teamxbox.com, which she knows to be high-quality sites for game reviews. The game titles from the inventory data could then be selected to drive that web search. If Ann had a real-time pricing and in-stock service available, it too could be included as service-based supplemental content.

2.2 Application Execution

Some of the key components in Symphony that come into play when processing a query are highlighted in Figure 2. We continue with the video game store example from Section 2.1.1, and assume that the owner, Ann, has already published her application to her web site, GamerQueen.

When a customer submits a query on the GamerQueen site, the auto-generated JavaScript on that page immediately forwards the query to Symphony for processing. The Symphony runtime must now process the query according to the configuration of the GamerQueen application. The query is first processed by the primary content sources, which in this scenario means searching GamerQueen's proprietary inventory data. In a more complex scenario, customer data could also be included to alter the query to, say, prefer some types of games over others.

Next, the supplemental content sources are queried based on data fields from the results from the primary content. The fields that should be used as arguments in these queries are specified by the application designer in the configuration file for the application.

In this scenario, a focused web search retrieves links to game reviews for each game result. This supplemental content is then merged with the selected display fields from the GamerQueen proprietary data and formatted into HTML, applying any configured layout and presentation details. The resulting HTML is then sent back to the embedded JavaScript, which then injects it into the GamerQueen page as a list of game results complete with reviews from the web.

Figure 3 shows an example of a search results page. In this case, a query "xbox" retrieves a number of games that can be played on an Xbox. Each result features fields retrieved from the inventory such as title (e.g., "Hail to the Chimp – Xbox 360"), image, and description. In addition, it also includes two hyperlinks (or any other specified number) to relevant game reviews retrieved from third-party review sites. Relevant text ads appear on the right-hand rail.

Note that to the customer, the entire experience appears as though it occurred seamlessly within the GamerQueen web site, even though all the processing and heavy-lifting was off-loaded to Symphony.

3. RELATED WORK

There have been prior attempts, mostly by commercial entities, to assist users in the creation of custom search engines. Most notably, Yahoo! BOSS is a SDK targeted at developers who wish to build search applications on top of Yahoo!'s search infrastructure. Its offering consists of an API to make query calls to Yahoo!'s search index, a client-side library to create custom mashups, and a possible partnership which allows specific third parties to tightly integrate their proprietary content with Yahoo!'s

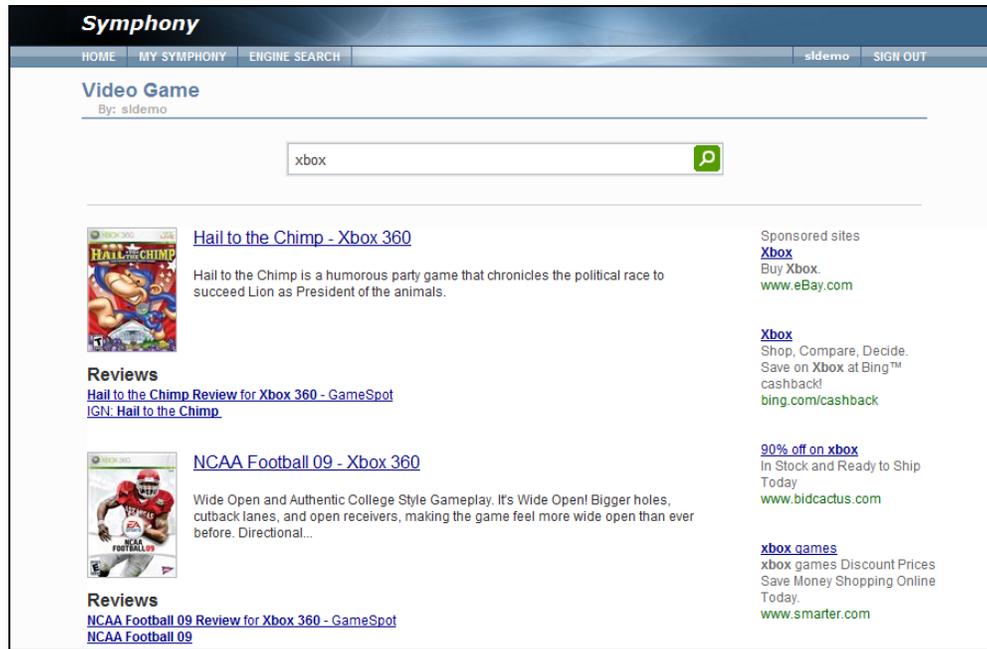


Figure 3: Search Results

infrastructure. Unlike Yahoo! BOSS, Symphony goes beyond simply opening up Live's search infrastructure to third parties. It aims at making the creation of a custom search engine as painless and seamless as possible to the average lay user, by providing additional support in the form of no-code user-friendly interface, hosting, and distribution of custom search engines. Yet other offerings such as those from Rollyo (rollyo.com) and Eurekster (eurekster.com) are even more limited than Yahoo! BOSS. They simply offer a site-restriction capability to focus searches on a restricted set of sites.

The capability to integrate information from several data sources to produce unified search results is somewhat similar to the notion of Web mashups [1]. Several commercial systems, such as Microsoft Popfly (popfly.com) and Yahoo! Pipes (pipes.yahoo.com), assist users in the creation of mashups by offering a programmingless interface to pull together several data sources. However, unlike Symphony, they do not offer access to integration with proprietary content. Neither do they offer indexing or searching over mashed up data. There are other proposals for making mashup creation easier, for instance with, spreadsheet [3], network navigation interface [4], or scripting language [5]. However, such implementations still require some form of learning effort on the part of the users (e.g., learning a scripting language).

4. CONCLUDING REMARKS

We have presented an overview of the Symphony platform and shown how it can be used to create custom search applications that leverage both a user's proprietary content and his or her domain expertise. As future work, we will explore how Symphony could further facilitate the creation of rich search applications through such means as recommending suitable supplemental

content (e.g., good game review sites) for a designer's primary content (e.g., game inventory), supporting richer querying of structured data, adding support for social search features, and creating new applications by composing other applications.

5. REFERENCES

- [1] Abiteboul, S., Greenshpan, O., and Milo, T. Modeling the mashup space. In *Proceedings of the 10th ACM International Workshop on Web Information and Data Management (WIDM'08)* (Napa Valley, California, USA, October 30, 2008). ACM Press, New York, NY, 2008, 87-94.
- [2] Fuxman, A., Tsaparas, P., Kannan, A., and Agrawal, R. Using the wisdom of the crowds for keyword generation. In *Proceedings of the International World Wide Web Conference (WWW'08)* (Beijing, China, April 21-25, 2008). ACM Press, New York, NY, 2008, 61-70.
- [3] Kongdenfha, W., Benatallah, B., Vayssière, J., Saint-Paul, R., and Casati, F. Rapid development of spreadsheet-based web mashups. In *Proceedings of the International World Wide Web Conference (WWW'09)* (Madrid, Spain, April 20-24, 2009). ACM Press, New York, NY, 2009, 851-860.
- [4] Lu, B., Wu, Z., Ni, Y., Xie, G., Zhou, C., and Chen, H. sMash: semantic-based mashup navigation for data api network. In *Proceedings of the International World Wide Web Conference (WWW'09)* (Madrid, Spain, April 20-24, 2009). ACM Press, New York, NY, 2009, 1133-1134.
- [5] Sabbouh, M., Higginson, J., Semy, S., and Gagne, D. Web mashup scripting language. In *Proceedings of the International World Wide Web Conference (WWW'07)* (Banff, Canada, May 8-12, 2007). ACM Press, New York, NY, 2007, 1305-1306.