

“Places” spam – the new front in the spam wars.

John Nagle
SiteTruth
nagle@sitetruth.com

December, 2010

Introduction

On October 27, 2010, Google released a major change in their primary search engine. For the first time, results from the “Google Places” system, previously confined to map-related searches, were merged into Google's main search results. Search results now contain more information about local businesses, and those search results appear prominently, near the top of web search results.

With the top results in Google searches now driven by the Google Places system, “search engine optimization” (SEO) efforts have been refocused on achieving a high rank in Google Places. This background paper discusses some of the results of those efforts.

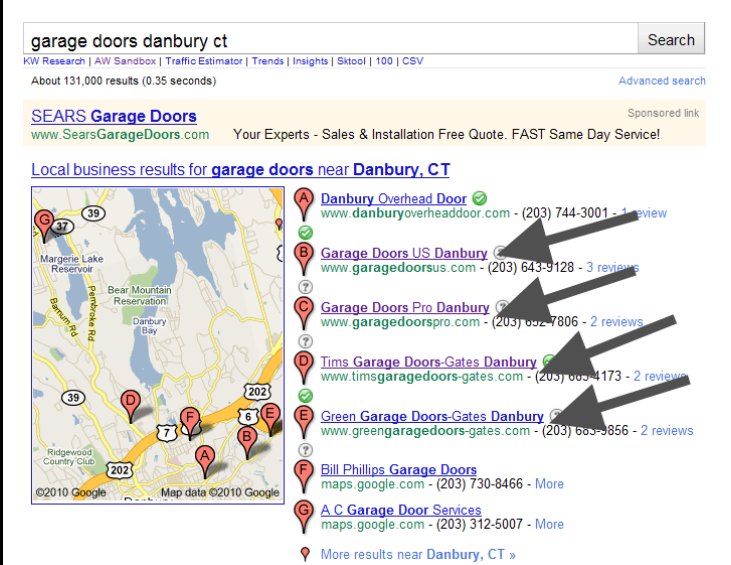
Possible solutions will be discussed in another paper. Here, the focus is on the scale of the problem.

“Search engine optimization” vs. Google Places

Until the October 2010 changes to Google web search results, the SEO community hadn't bothered to heavily target Google Places or Google Maps. That has changed. In only one month, the SEO industry developed techniques not only to achieve high positions in Google Places, but to create totally phony place entries.

On the right is a Google Maps entry showing four completely phony entries. As of today, a Google web search for “garage doors danbury ct” returns three Google Places results, the second of which is from that set of phony entries. All the phony entries still appear in Google Places.

For this particular set of phony results, we know how it was done,



The screenshot shows a Google search for "garage doors danbury ct". The search results include a sponsored link for SEARS Garage Doors and a list of local business results for "garage doors near Danbury, CT". The map shows several phony entries marked with letters A through G. Arrows point from the text below to these entries.

Those businesses don't exist.
The tricks that put them there are described in "Dominating Google Maps- The Most Effective Spam Ever And What You Can Learn From It"

because the search optimization company which did it, “Convert Offline”, uses it as an example to market their services. They describe what they did:

1. *Set up a listing in Google Maps at an address that does not currently exist. For example, where there is a 60 Main St., Anytown and a 64 Main St, Anytown and these represent real addresses. Set up your listing at 62 Main St.*
2. *Name your business USKeyword-City, or Keyword-Pro-city or Fictitious name of person plus keyword for the personal touch.*
3. *Build citations to your listing. These listings contain citations from Yahoo Local, Hotfrog, Guidespot, local.newstimeline.com.*
4. *Create a blog on one of the sites for the purpose of creating a perfect citation for thousands of listings.*
5. *Link build*
6. *Give your new listing a sparkling review*
7. *Now find an adjacent town and repeat. Again and again and again again.*

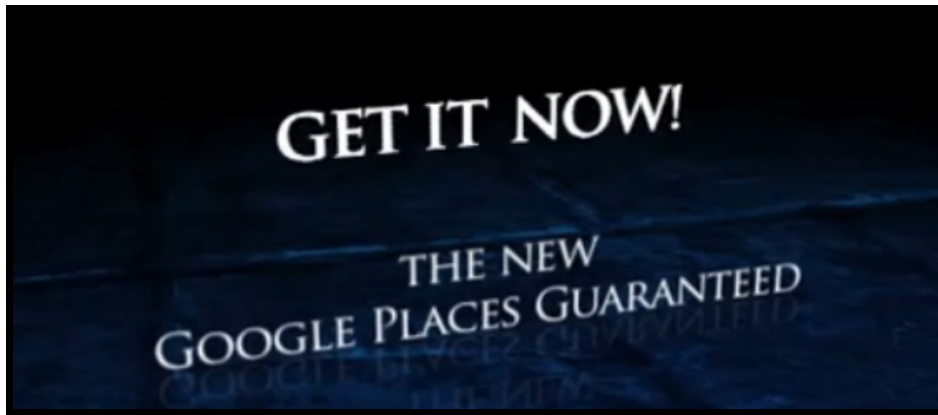
Similar examples are easy to find. Google's results for “New York City locksmith” reveal a similar problem. Google's own forums for customer complaints about Google Pages reveal other areas under assault. Google lacks an effective strategy for dealing with this form of attack.

Industrial-strength spam

Google Places has been integrated with Google web search for only two months. The SEO community already has services in place to spam it. The “white hat” search engine optimization industry, which generally tries to avoid committing felonies, is sufficiently confident of their ability to push Google Places listings upward in search results to offer performance guarantees.

The techniques used typically involve generating large numbers of phony recommendations. The companies offering these services are quite open about admitting that they do this. This has perhaps become an accepted promotional technique, and is destroying the value of recommendation sites.

The line between “legitimate” and “black hat” search engine optimization has been blurred in the Google Places arena. Because spamming Google Places is both easy and competitively essential, the legitimate players have entered the fray. If allowed to establish themselves, they may be difficult to stop in future.



Musson Media

Top placement guaranteed or your money back



SEO Expert Global

Rank at the top of Google Places. "100% Satisfaction Guaranteed"

Google Places SEO Packages

Google Places SEO Packages	Paper Boy Package	Courier Package	The Postal Service Package	Freight Shipping Package
Stay Ranked For: 3-4 months	This package lives up to its name. Hit your reputation on Google Places full force with 100 stellar reviews from your happy customers and watch your lead stream grow virally.			Paypal


365Outsource

*"100 stellar reviews from your happy customers."
Includes "user farming resourcing" and
"Geo Located Smart IP service"*

Ads from "legitimate" search engine optimization firms promising top placement in Google Places results.

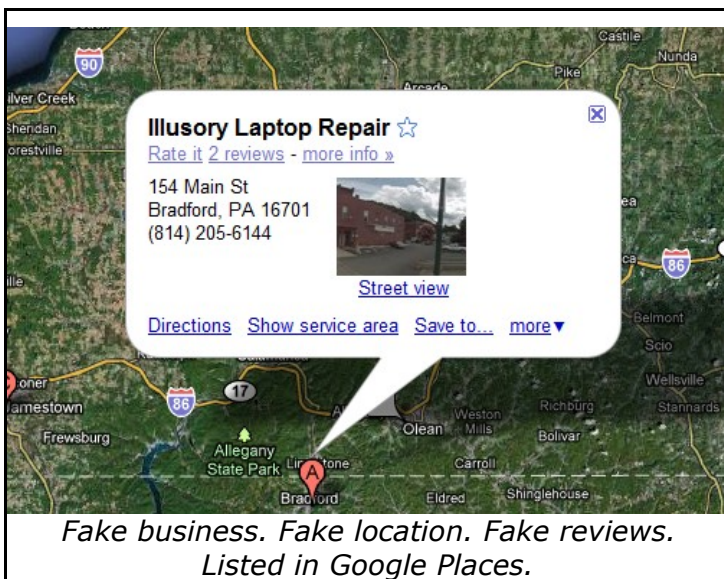
New “Black Hat” Techniques for Places

On the “black hat” side of the world, where techniques verge on criminal activity, more aggressive techniques are in use. Some are effective; others are just amusing.



Some commercial products used for attacking Google Places

"Convert Offline", as mentioned previously, advertises "Dominating Google Maps - The Most Effective Spam Ever And What You Can Learn From It". They recommend inserting phony Google Places entries using phony business addresses on real streets.



A more elaborate scheme is "Basics Plus New York Directory". This is a setup by a company that has a chain of convenience stores in New York. They let other companies use their locations as semi-bogus addresses. They have a deal with "Paragon Locksmith" which makes that company appear to be all over town. This trick propelled that locksmith company into the Google Places top result pack, and totally dominated Bing's places listing.

Blumenthals, a small company in the web marketing business, deliberately inserted a completely phony business, "Illusory Laptop Repair", into Google Places, as a demonstration of the weakness of Google's defenses. They


discovered a way to bypass Google's postcard verification system.

Completely phony addresses for businesses are quite common. This is most common for businesses which are very densely located in urban areas, provide their services off-site, and wish to appear to have more physical locations. A few categories are already dominated by such phony addresses:

- Locksmith
- Plumber
- Carpet cleaning
- Movers
- Appliance Repair

To date, most of the Google Places spamming operations seem to involve substantial human effort. Some of that effort is outsourced to low-wage

countries. The “black hat” community is working on high-speed software-implemented bulk Google Places spamming tools. From discussions on “black hat” forums and ads for developers, it appears that several parties are trying to adapt the Xrumer blog spam engine, which automatically inserts blog and forum spam into a large number of sites, to generate large numbers of phony recommendations automatically.



Xrumer

*"Dominate Google and
Crush Your Competition –
Get Your #1 Spot Now"*

*From Botmaster Labs,
Kiev, Ukraine*

There are mentions on “black hat” forums of a new spam engine, specifically targeted at Google Places. to be released in January 2011. Once this process is fully automated, the volume of phony entries in Google Places can be expected to increase substantially.

How Craigslist lost its spam war

A similar battle was fought on Craigslist. Spam on Craigslist had been a minor nuisance since the early days of Craigslist. That changed two years ago. In 2008, the spammers started winning and took over much of Craigslist. Commercial firms openly advertised products and services for of spamming Craigslist.

Craigslist tried to stop spamming by checking for duplicate submissions. They checked for excessive posts from a single IP address. They required users to register with a valid E-mail address. They added a CAPTCHA to stop automated posting tools. And users could flag postings they recognize as spam. Craigslist had in place all the standard best practices.

Those measures failed. Several commercial products were developed to overcome obstacles to bulk posting. CL Auto Posting Tool is one such product. It not only posts to Craigslist automatically, it has built-in strategies to overcome each Craigslist anti-spam mechanism.

Spammers added random text automatically to each spam message to fool Craigslist's duplicate message detector. IP proxy sites were used to post from a wide range of IP addresses. Those attacks were effective, and Craigslist was unable to counter them.

Craigslist, like Google Places, tried E-mail authentication, requiring a unique E-mail account for each Craigslist posting account. The spammers retaliated by using Jiffy Gmail Creator ("Who Else Wants to Create Unlimited Gmail Accounts in Seconds Flat Without Breaking a Sweat?") to create tens of thousands of dummy E-mail accounts.

Craigslist tried CAPTCHAs, letter puzzles which, supposedly, only humans could solve. That barely slowed the spammers down. A combination of OCR and outsourcing to low-wage countries overcame that defense.

Craigslist tried manual flagging of spam. That didn't work either. CL Auto Poster has an automatic monitoring system which detects when a posting has been flagged as spam and re-posts it.

CL Auto Poster isn't the only such tool. Other desktop software products are AdBomber and Ad Master. For spammers preferring a service-oriented approach, there's ItsYourPost.



With these power tools, the defenses of Craigslist were overrun. Some categories on Craigslist became over 90% spam. The personals sections were the first to go, then the services categories, and more recently, the job postings.

Finally, Craigslist tried verifying users by telephone. Posting in some categories required a callback phone call, with a password sent to the user either by voice or as an SMS message. Only one account is allowed per phone number. Only these “phone verified accounts” (PVA) were allowed to post in key categories. Spammers reacted by using free VoIP numbers. Craigslist blocked those. Spammers tried using number-portability services such as Grand Central and Tossable Digits. Craigslist blocked those. Spammers tried using their own free ringtone sites to get users to accept the Craigslist verification call, then type in the password from the voice message.



Commercial services arose which obtained access to large blocks of phone numbers obtained from grey-market phone companies, and could provide large numbers of seemingly valid phone numbers on a short term basis. This provided a unique phone number for each phony posting account.

Most of the defensive technique currently being used by Google Places were tried by Craigslist, and failed. The spammers won.

The fundamental vulnerabilities in Google Places

The two phases of spamming Google Places are the insertion of fake business locations and the creation of fake reviews. Both are embarrassingly easy using the techniques described above.

Google Places obtains business locations from web pages created by the business itself, advertising directories (“Yellow Pages”) entries paid for by the business, and from “place pages”, also created by the business itself. There is little if any verification against objective data sources, such as business licenses, corporation registrations, and business credit rating services such as Dun and Bradstreet. This makes it possible to create fake Google Places entries.

Recommendations are obtained from recommendation web sites. Most recommendation sites allow free account creation and have little information about their members, so the cost of creating phony identities for recommendation spam is low. Because the typical local business has a relatively small number of recommendations, only a few phony recommendations are needed to promote an individual business location.

Spamming recommendation sites is cheap. Fake recommendations and reviews can be created inexpensively and with low risk. The comparable attack for organic search, creating “link farms” of junk sites and fake blogs linking to more junk sites, is more difficult and costly. Link farming involves hosting and site maintenance. Link farms which are identified as such by search engines may suddenly lose their value, destroying the spammer's investment. Recommendation spam does not carry that financial risk, since innocent third parties host the fake recommendations for free. Spamming Google Places is thus much cheaper than

spamming organic search, and is a likely growth area for aggressive spam operations.

Industry commentary

Influential trade sites have started to pick up that something has gone wrong at Google. All the items below are from December 2010.

- **“Has Google Jumped the Shark”**
(Kristine Schachinger, Search Engine Watch, December 16, 2010):

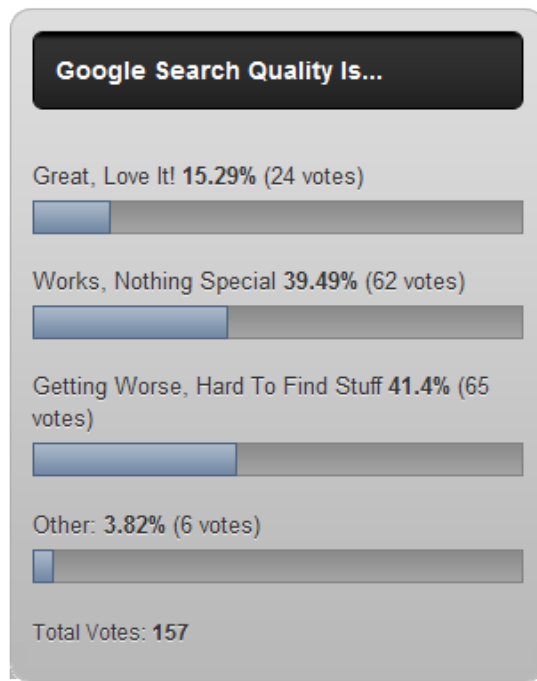
*“Oh, Search! Is That What We Do? There has just **been so much wrong lately** -- results filled with top 10 AdSense sites, 404 pages, irrelevant terms, lack of real authority sites, big brand domination, and the list goes on and on. ... What I don't get are good results for my search entry. Why? ... When I was at Pubcon, Matt Cutts mentioned that the **Google engineers responsible for spam control and other organic issues had been spending a lot of time outside of their normal focus, the algo.** (Guess they thought we wouldn't notice?)”*

- **“Google’s SMB Achilles’ Heel: People”**
(Frank Reed, Marketing Pilgrim, December 24, 2010):

*“I don’t believe that they realize the issues that exist among many of the 4 million or so verified place page owners. As a result, **what they think is a sales call will very likely turn into questions and a request for service** that Google will have to answer and support before they can “close a sale”. This request for service will reveal Google’s true weak spot: it doesn’t understand people. ... Visit the Google Places forum to get a taste of what is **wrong with the Place Page system.** Duplicate listings, confusing processes, lack of human interaction, misplaced reviews, listings coming and going with no real reason ... But considering how important these pages appear to be to Google it is simply mind blowing that they think that they work just fine as they are and they don’t need to truly support the businesses that they are trying to extract more money from through this offering. **Google is asking for trouble here and I hope they get a lot of it.**”*

- **“Google Places Integrated Results Are Less Than Reliable”**
(Jennifer Eaton, Standing Dog Marketing, December 15, 2010)

“I have long maintained that if Google wants to be in the business of being the world’s Yellow Pages – which is what this is doing – accuracy must be taken to the next level.



Search Engine Roundtable poll
October - December 2010

*... Once Google starts saying “Here is where business A is located and their telephone number,” **there is very much a right and wrong answer. When they get it wrong, they cost businesses money.**”*

- **“Google Map Spam Creeping into the Hinterlands”**
(Jill Whelan, Sphinn Internet Marketing Forums, December 8, 2010)

*“By making Places Pages so prominently featured in search results right now, **Google has opened themselves up to mega super duper spamming.** For some, it's going to be the only way to compete with Places Pages.”*

Conclusion

Spamming Google search results is easier and cheaper since the merger of Google Places results into web search. In only two months, effective techniques for spamming Google Places have come into wide use. Search quality as perceived by users is deteriorating. Industry sources are critical of Google's inability to deal with the problem.