A Walk Up the Stack

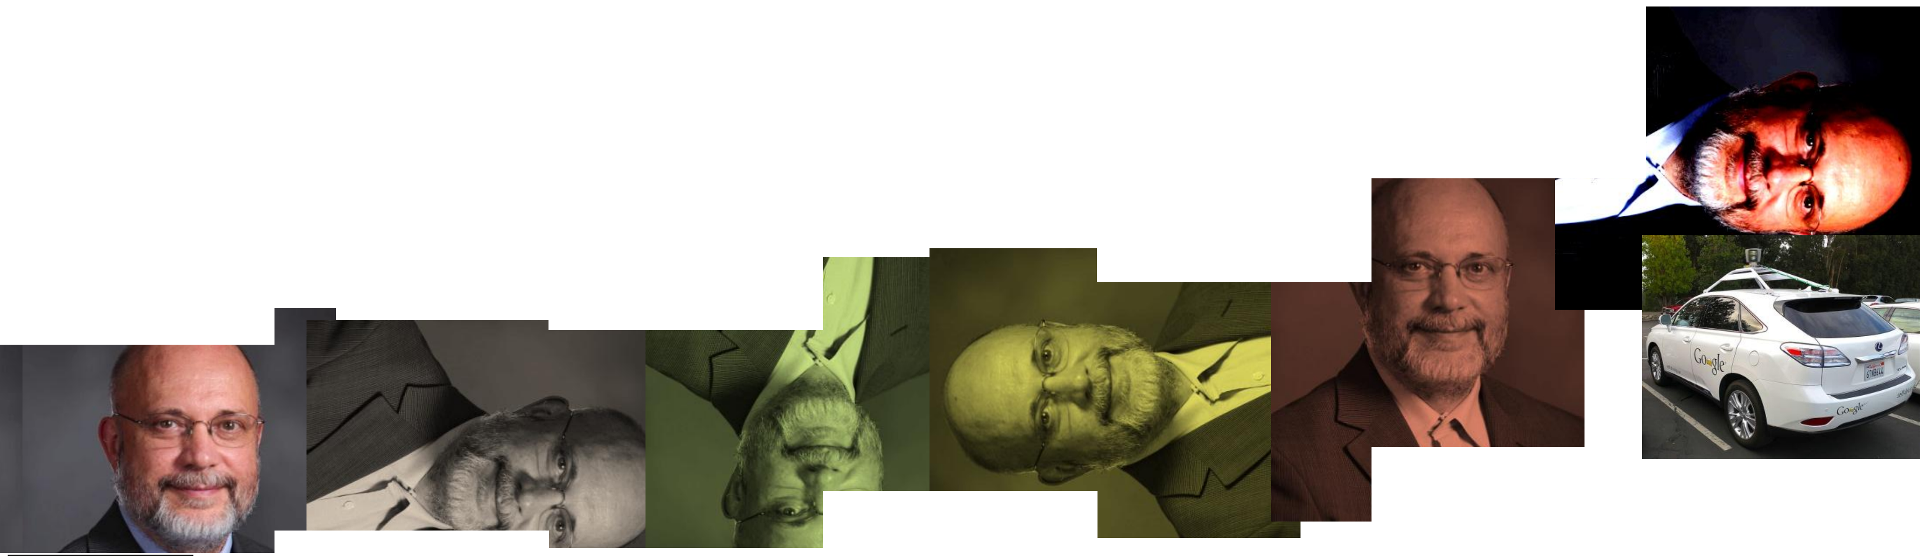

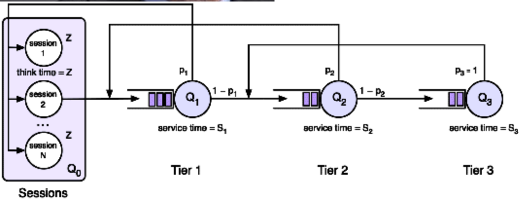

$$\bigoplus_{i=0}^{T-1} \left(A^{\otimes i} \right.$$

Jean Bolot
technicolor

1. Intersections with François

First paper read

An End-to-End Approach to the Resequencing Problem

FRANÇOIS BACCELLI

INRIA, Le Chesnay, France

AND

EROL GELENBE AND BRIGITTE PLATEAU

Université de Paris-Sud, Orsay, France

Abstract. The resequencing or serialization problem is of basic interest in distributed systems and computer communication systems. This is because a flow of packets, messages, or updates entering a communication system in chronological order from the same port or from different ports may be disordered. The receiving port must then ensure that these objects are resequenced in the appropriate order before they are fed to the output of the system. In this paper we analyze the end-to-end delay incurred by objects traversing such a system, including the disordering delay, the delay introduced by the resequencing algorithm, and the delay due to the output server at the receiving port. The analysis is carried out via factorization methods.

Categories and Subject Descriptors: D.4.8 [**Operating Systems**]: Performance—*queuing theory*

General Terms: Performance, Theory

Additional Key Words and Phrases: Consistency control, packet-switching networks, performance analysis, serialization

1. *Introduction*

The resequencing problem is a fundamental issue in networks and in distributed systems. Let us first give an abstract statement and then provide examples of some practical occurrences of the problem.

Consider a sequence of objects $\{\phi_n\}_{n\in\mathbb{N}}$, where $\mathbb{N}$ denotes the set of all nonnegative integers. They enter a communciation system at instants $\{a_n\}_{n\in\mathbb{N}}$ where a_n corresponds to ϕ_n.

Each ϕ_n is then delayed by some time D_n, $n \in \mathbb{N}$. Thus at the output point of the system, the objects appear at instants $\{a_n + D_n\}_{n\in\mathbb{N}}$, but these instants are not necessarily in chronological order any more (i.e., it is possible that $a_n + D_n > a_l + D_l$ for $l > n$).

These objects are then processed by the resequencing algorithm (RA); ϕ_n will receive some service of duration S_n and depart at time d_n. However this service can only be given in the same order as that of the external arrival instants; that is, ϕ_n

Authors addresses: F. Baccelli, INRIA, Domaine de Voluceau, BP 105-Rocquencourt, 78153 Le Chesnay Cedex, France; E. Gelenbe and B. Plateau, ISEM-Bât. 508, Université de Paris-Sud, 91405 Orsay, France.

Journal of the Association for Computing Machinery, Vol. 31, No 3, July 1984, pp. 474–485.

First paper co-authored

Discrete-Time Analysis of Adaptive Rate Control Mechanisms

Eitan Altman[1], François Baccelli[1], and Jean-Chrysostome Bolot
INRIA, B.P. 93, 06902 Sophia Antipolis Cedex, France

Abstract

We analyze the performance of a generic feedback flow control mechanism which captures the properties of several such mechanisms recently proposed in the literature. These mechanisms dynamically regulate the rate of data flow into a network based on feedback information about the network state. They are used in a variety of networks and they have been advocated for upcoming high-speed networks. However, they are difficult to model realistically. In this paper, we present a stochastic discrete-time approach that yields models which are realistic and yet tractable and computationally easy to solve.

For our generic mechanism, the feedback consists of an exponentially averaged estimate of the bottleneck service rate and queue size. We obtain a model described by non-linear stochastic difference equations. We find the conditions under which these equations converge to a steady-state and we characterize the speed of convergence to steady-state. We then consider a linearized version of the model for which we can derive closed-form solutions. These solutions bring out a tradeoff between efficient steady-state behavior and rapid adaptability to varying network conditions. We indicate how a dynamic averaging technique of the feedback information can overcome this disadvantage. We also identify the cases when the linear model is a good approximation to the non-linear model.

Keyword Codes: C.2.0; C.2.1; C.4
Keywords: Computer-Communication Networks, General; Network Architecture and Design; Performance of Systems

1 Introduction

In a packet-switched network, packets generated by source nodes are routed via a sequence of intermediate nodes to destination nodes. A flow control mechanism limits source rates in order to avoid congestion in the intermediate nodes. The goal of the flow control mechanism is to match the source rate on a connection to the capacity available on this

[1]Supported by the European Grant BRA-QMIPS of CEC DG XIII.

Last paper co-authored

Modeling the Economic Value of Location and Preference Data of Mobile Users

F. Baccelli (ENS/INRIA, francois.baccelli@ens.fr) and J. Bolot (Sprint Labs, jeanbolot@gmail.com)

Abstract—The defining characteristic of wireless and mobile networking is user mobility, and related to it is the ability for the network to capture (at least partial) information on where users are located and how users change location over time. Information about location is becoming critical, and therefore valuable, for an increasingly larger number of location-based or location-aware services. One key open question, however, is how valuable exactly this information is.

Our goal in this paper is to develop an analytic framework, namely models and the techniques to solve them, to help quantify the economics of location information. Our aim is to derive models which can be used as decision making tools for entities interested in or involved in the location data economics chain, such as mobile operators or providers of location aware services (mobile advertising, etc). We consider in particular the fundamental problem of quantifying the value of different granularities of location information, for example how much more valuable is it to know the GPS location of a mobile user compared to only knowing the access point, or the cell tower, that the user is associated with. We illustrate our approach by considering what is arguably the quintessential location-based service, namely proximity-based advertising.

We make three main contributions. First, we develop several novel models, based on stochastic geometry, which capture the location-based economic activity of mobile users with diverse sets of preferences or interests. Second, we derive closed-form analytic solutions for the economic value generated by those users. Third, we augment the models to consider uncertainty about the users' location, and derive expressions for the economic value generated with different granularities of location information.

To our knowledge, this paper is the first one to present and analyze economic models which can help understand the economic value generated by mobile users with location based services, for different granularities of location information in wireless networks.

I. INTRODUCTION

dial up and broadband access fees. Furthermore, the number of users of the mobile Internet (measured by the number of users accessing browser-based services on cell phones only) is estimated at between 500 million and 1 billion, almost on par with the total number of PCs connected to the Internet [1], [2].

A key characteristic of mobile networks and devices is their ability to capture and analyze (at least partial) information on the location of mobile users. For example, cellular operators have routinely captured large scale location data for billing purposes, but also to improve location management or satisfy legal requirements such as E911. More recently, they have started exposing (often for a fee) large scale location information to application developers. In parallel, a significant fraction of mobile devices is now GPS-enabled and captures and (sometimes for a fee) provides access to real-time GPS data. Note that the capture and availability of location data is part of a larger trend, where data of various kinds such as location data but also social network data or spectrum usage data is seen as an extremely valuable, even strategically important, asset by the carriers in particular, but also by the general mobile industry.

In any case, the capture and availability of location data enables the development of a wide range of location-based or location-aware services, and indeed an rapidly increasing number of such services is now available, ranging from navigation to location-aware advertising, friend finder, etc, and many more are announced or launched on a daily basis. As noted above, this location data, since it enables new services and new economic activities, is also seen as economically valuable. This raises the question then of how valuable it is,

Most used paper

Bayesian Inference for Localization in Cellular Networks

Hui Zang
Sprint, USA

Francois Baccelli
ENS, France

Jean Bolot
Sprint, USA

Abstract—**In this paper, we present a general technique based on Bayesian inference to locate mobiles in cellular networks. We study the problem of localizing users in a cellular network for calls with information regarding only one base station and hence triangulation or trilateration cannot be performed. In our call data records, this happens more than 50% of time. We show how to localize mobiles based on our knowledge of the network layout and how to incorporate additional information such as round-trip-time and signal to noise and interference ratio (SINR) measurements. We study important parameters used in this Bayesian method through mining call data records and matching GPS records and obtain their distribution or typical values. We validate our localization technique in a commercial network with a few thousand emergency calls. The results show that the Bayesian method can reduce the localization error by 20% compared to a blind approach and the accuracy of localization can be further improved by refining the *a priori* user distribution in the Bayesian technique.**

I. Introduction

Mobile phones have become a fundamental component of modern lives and economies, and have become ubiquitous, reaching an estimated 4.1 billion by end of 2008 - over half of the planet's population - with several countries having penetration rates much higher than 100% [1]. As a result, recent research is looking at mobile phones as a powerful and exciting new tool to track and analyze human behavior, in particular human social interactions and activity patterns [2], [3].

Understanding human mobility patterns is of major importance in a number of areas, including of course cellular network design and engineering, but also urban planning, transportation geography, (mobile) advertising, crowd and event management, or epidemics monitoring and control. Until relatively recently, few tools and little large scale data were available to monitor the spatial dynamics of large populations of users. (For example, refer to [4] for a study of human mobility and travel patterns by studying the circulation of bank notes in the United States). The call data records (CDRs) collected by wireless operators for billing and troubleshooting purposes now provide one such source of data, and they make it possible to study human mobility patterns of populations at previously impossible-to-achieve scales. The main challenge with CDRs has been availability. It remains limited, in particular because of privacy concerns and corresponding challenges with anonymization [5], but still availability has become greater in the past couple of years. We anticipate that there will be more work on CDRs and human mobility modeling can greatly benefit from the increasing availability of CDRs.

Our work is motivated by the need to convert large sets of CDRs to location records. CDRs contain only coarse-grained location information about the mobile user as cell and sector IDs and round-trip-time (RTT) and signal-to-noise-and-interference ratio (SINR) measurements. We would like to obtain more fine-grained location information. Although there has been a lot of work on localizing mobiles in a cellular network, most of them are based on the operation mode in which a mobile could see a few surrounding base stations [6], [7], [8], including the ones with very weak signals. However, only base stations with strong enough signals can carry a mobile's communication and only the base stations that were actually carrying the mobiles' communication were logged into CDRs. Although during handoffs, would two or more base stations carry a call simultaneously, in all other cases, 50% of our CDRs, the mobile is served by *only one* base station. In order to convert CDRs to location records, we need to deal with the majority cases where there is only one base station recorded. Clearly popular localization techniques such as TOA, AOA [6], [9], or signal-strength fingerprinting [10] are not applicable because they require two or more distinct base stations. Therefore, we need to solve the problem of localization for these single-legged calls.

We develop a Bayesian-based method to localize users in cellular network with only one call-leg information. We consider information such as the distance to the base station, location of neighboring base stations, and levels of interference and/or noise. We demonstrate the benefits of our technique with CDRs of E911 calls with matching GPS coordinates. We can improve the localization accuracy by 20% comparing to a blind approach in which a location is randomly chosen along an arc of the sector and the arc is determined based on distribution of RTT measurements.

Although our work is motivated by *offline* call data processing, the technique developed in this work also benefits location-based mobile applications which do not have direct access to information from several cell/sectors to perform localization using the alternative methods (TOA, AOA, fingerprinting). Our approach is very general and applicable not just to cellular networks, but to other wireless networks in particular wireless LANs.

The rest of the paper is organized as follows. Section II provides background information about the network and the data set under study and review related work. We develop the Bayesian-based method in steps in Section III. Section IV presents measurement results that aid us in selecting parameters for the Bayesian method. The method is evaluated in Section V and Section VI concludes the paper.

II. Background

A. Network Information

We consider a commercial CDMA2000 network which carries voice, data and SMS traffic. We obtain a network map with locations of base stations as (latitude, longitude) pairs. In the network, all base stations are equipped with directional antennas and each cell has two or three sectors. We know the azimuth (direction) of each antenna, which corresponds to the

I-thought-would-be-most-used paper

The Role of PASTA in Network Measurement

François Baccelli, Sridhar Machiraju, *Member, IEEE*, Darryl Veitch, *Senior Member, IEEE*, and Jean Bolot

***Abstract*—Poisson Arrivals See Time Averages (PASTA) is a well-known property applicable to many stochastic systems. In active probing, PASTA is invoked to justify the sending of probe packets (or trains) at Poisson times in a variety of contexts. However, due to the diversity of aims and analysis techniques used in active probing, the benefits of Poisson-based measurement, and the utility and role of PASTA, are unclear. Using a combination of rigorous results and carefully constructed examples and counterexamples, we map out the issues involved and argue that PASTA is of very limited use in active probing. In particular, Poisson probes are not unique in their ability to sample without bias. Furthermore, PASTA ignores the issue of estimation variance and the central need for an inversion phase to estimate the quantity of interest based on what is directly observable. We give concrete examples of when Poisson probes should not be used, explain why, and offer initial guidelines on suitable alternative sending processes.**

***Index Terms*—Active probing, network measurement, Nonintrusive Mixing Arrivals See Time Averages (NIMASTA), Poisson Arrivals See Time Averages (PASTA).**

I. Introduction

POISSON Arrivals See Time Averages, or "PASTA," is a property applicable to many stochastic systems. In essence, it states that observations made of a system at time instants obeying a Poisson process, when averaged, converge to give the 'true' value, that is, to the average that an ideal observer would make when monitoring the system continuously over time. PASTA was first formalized by probabilists, notably in the 1970s. Wolff, in his classic 1982 paper [24], unified and extended the then-existing PASTA results. The generality of his formulation, based on the "Lack of Anticipation Assumption" (LAA), which requires simply that the past history of the system does not influence the arrival times of future observers, did away with the need to prove ergodic theorems for each new application and led to PASTA being widely used.

PASTA has been used [15], [16], [22], [25] to justify the sending of probes (or probe trains) at Poisson epochs in an effort to obtain unbiased estimates of quantities of interest, for example, end-to-end *delay*. However, despite the generality of the PASTA result of Wolff, in many respects the role and utility of PASTA for active probing has become unclear both in the theoretical and practical senses. This paper aims to clarify what Poisson probing, and PASTA itself, can and cannot provide for active probing. In this context, key questions include:

- When is PASTA valid in the strict sense?
- When and in what sense is PASTA useful when it holds? Is Poisson probing necessarily optimal?
- Are there cases when Poisson probes should *not* be used?
- What role is played by PASTA within the inference problems of active probing?

Related to this last point, there is an important, prior question: **What does PASTA apply to**? In other words, Poisson arrivals see time averages, but of what? Does PASTA hold for *any* quantity that may form the object of active probing?

Our main focus in this paper is on end-to-end delay over a tandem queueing network, to which PASTA can in fact apply. Delay is a simple, yet important target of active probing measurement in its own right. A natural aim in this context would be to accurately determine any desired statistic of the delay that would be experienced by a single packet of any given size sent into the network in its steady state regime, for example, the distribution of such a delay. A particular case is the *virtual work* of queueing theory, which corresponds to the delay a zero-sized packet would see under FIFO scheduling when sent into the network in steady state. By carefully distinguishing between the nonintrusive case (virtual probes of zero size) and the intrusive case (real probes of finite size), we provide important insights into the above questions. The simplicity of delay allows rigorous results to be derived, and yet it provides a context rich enough to inform active probing techniques in general.

Our findings group naturally under three distinct categories and can be summarized as follows.

Sampling Bias versus Intrusiveness

- PASTA states that Poisson sampling is unbiased. In the nonintrusive case, we show that this is not unique to Poisson, but is shared by a large class of other sampling processes.
- PASTA states that Poisson sampling remains unbiased even when observers are not virtual but contribute to system load. Apart from a few exceptions ([13]), this

Manuscript received November 20, 2006; revised September 08, 2008; approved by IEEE/ACM Transactions on Networking Editor M. Roughan. First published July 14, 2009; current version published August 19, 2009.

F. Baccelli is with INRIA-ENS, Ecole Normale Supérieure, Paris 75230,

Move "Up the Stack"

From networks

To usage

Move "Up the Stack"

From networks

To usage

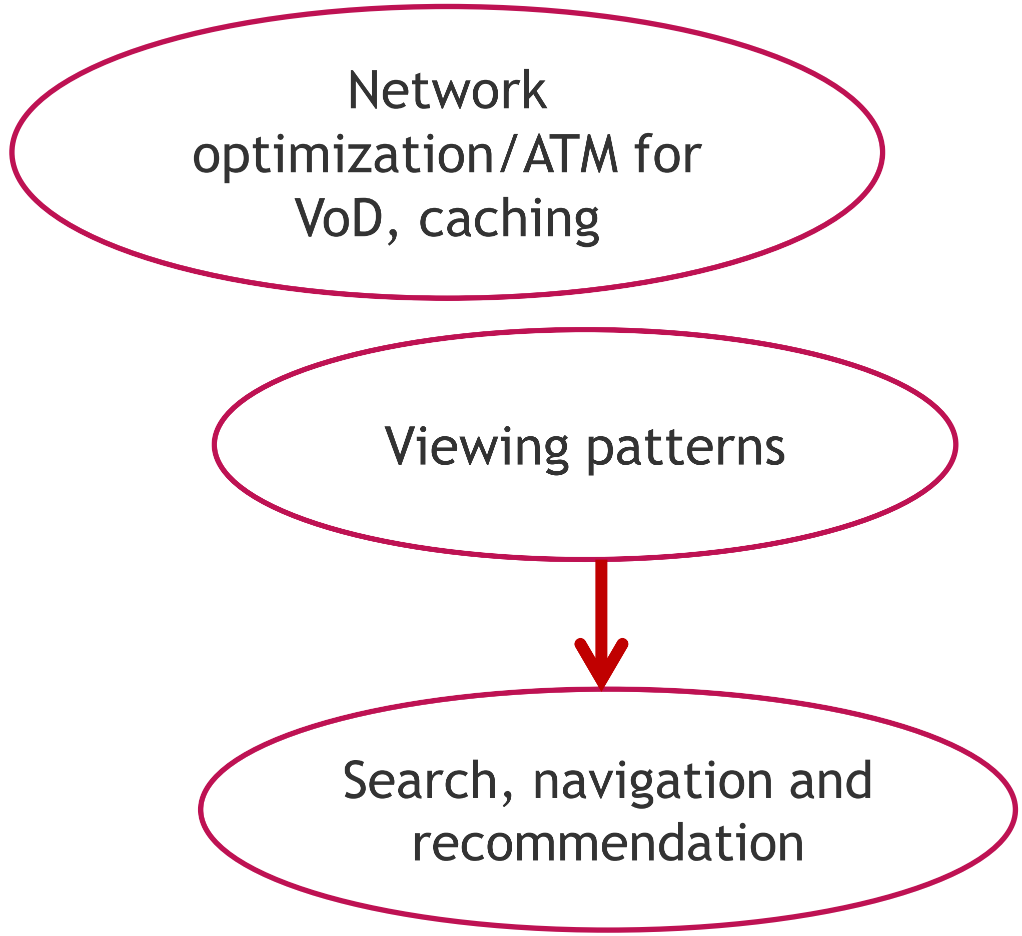

Move “Up the Stack”

From networks

To usage

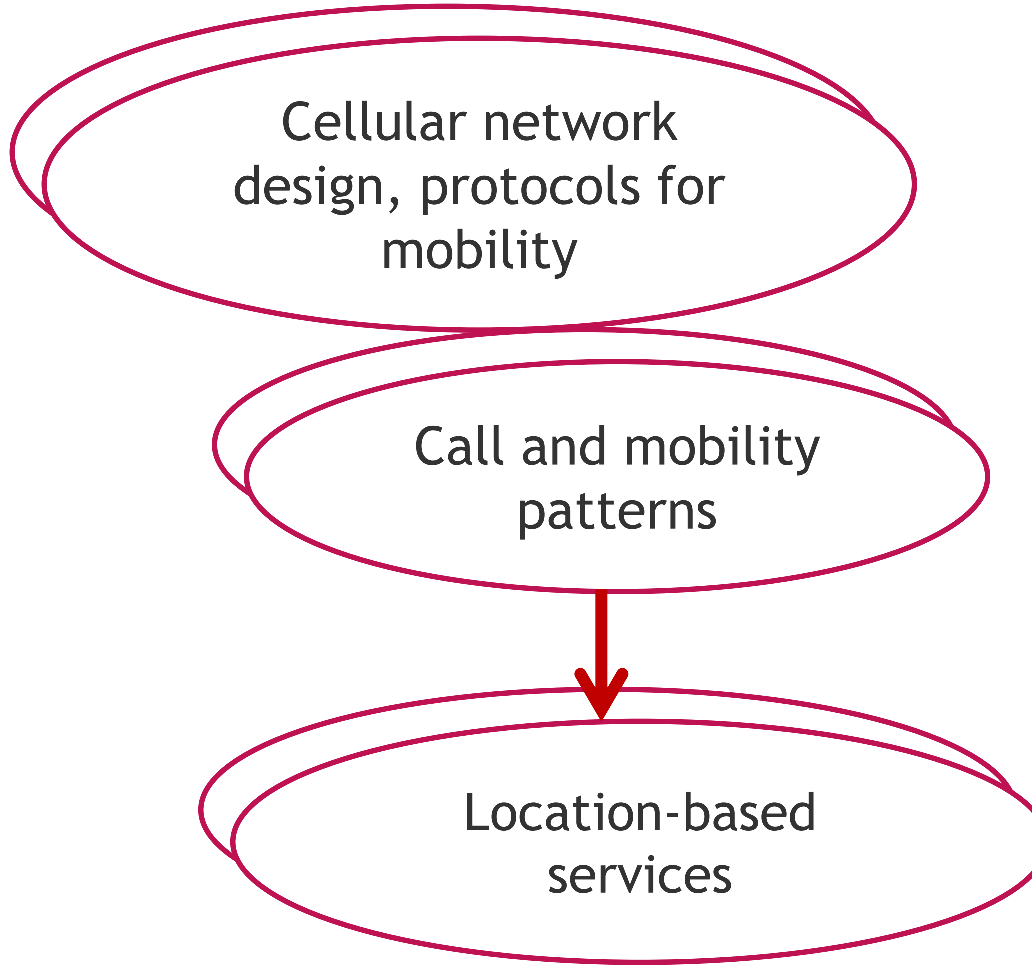

2. Modeling “Up the stack”

Stochastic processes and stochastic geometry just as important up the stack as they have been down the stack…

Stochastic geometry for: network design
location data

Quantify value of user location data

Exact

Approximate

Quantify value of user location data

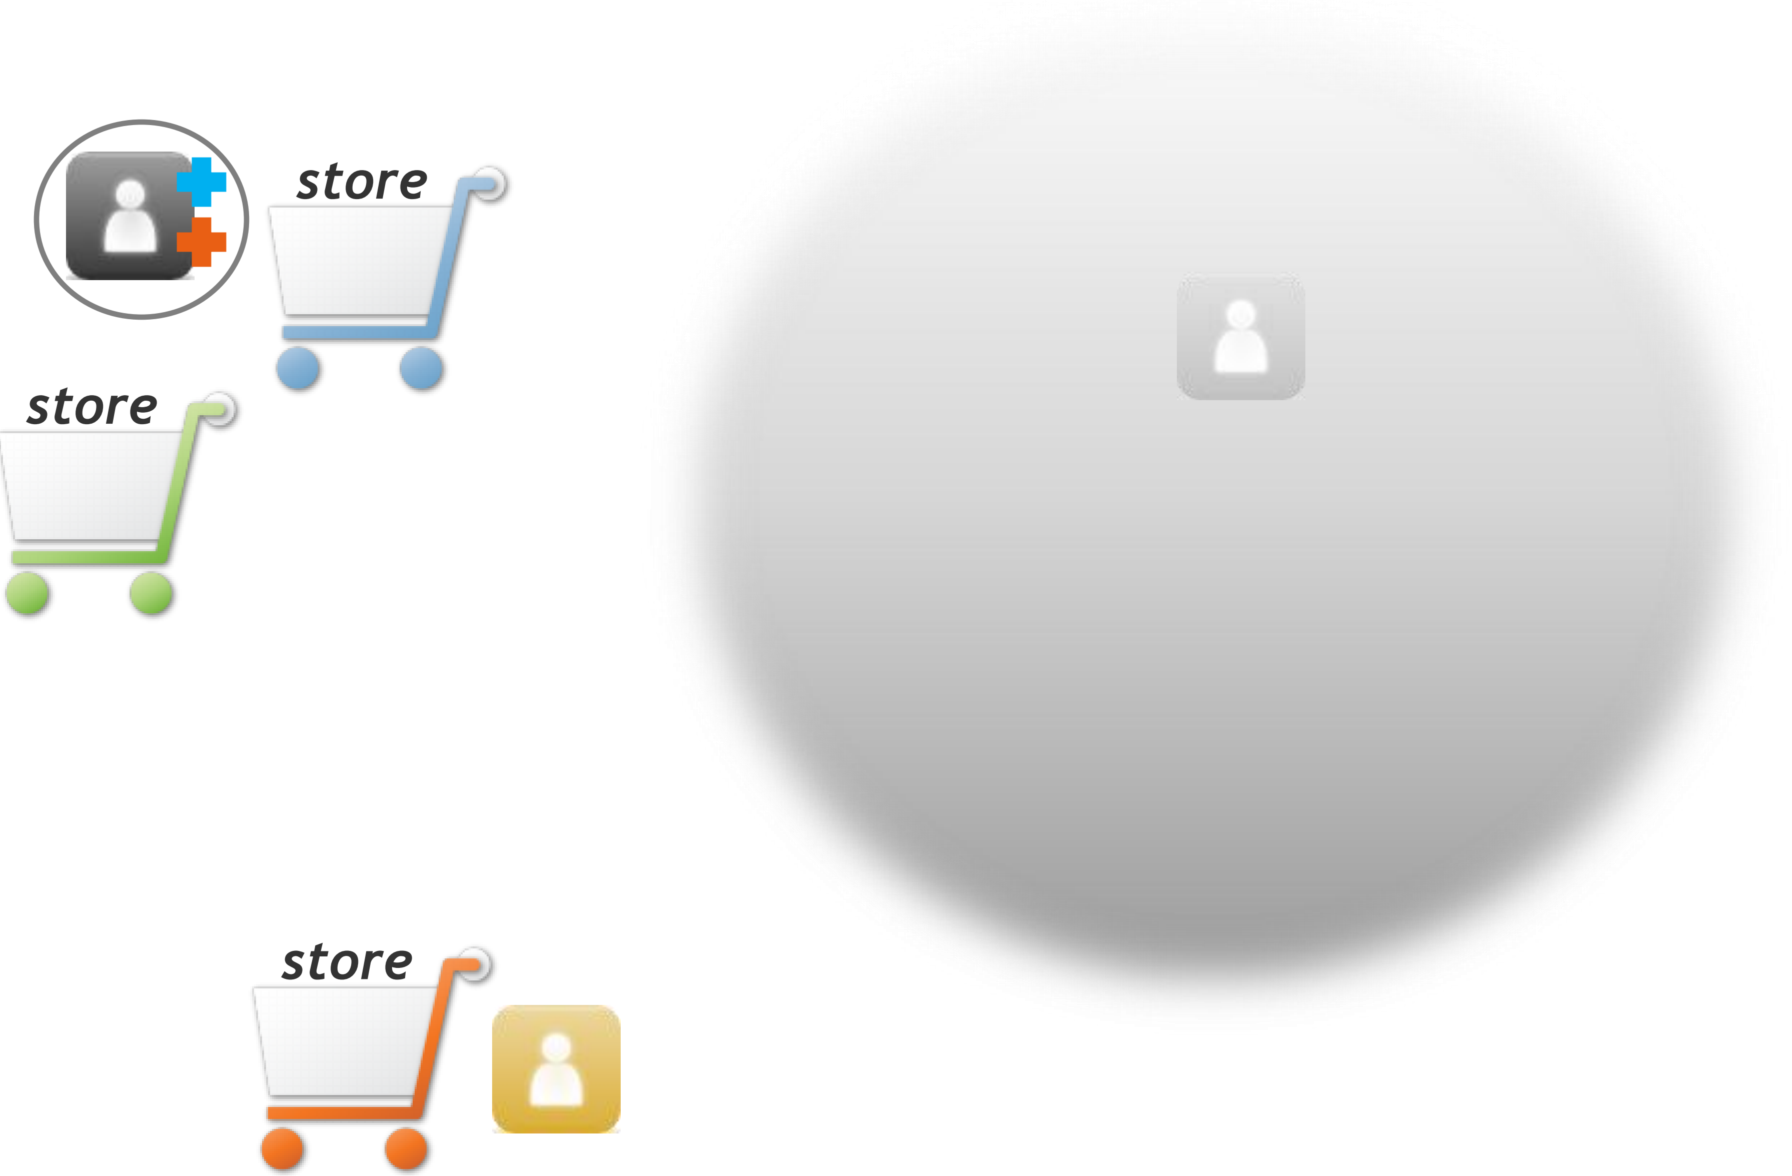

Approach: Location-based services

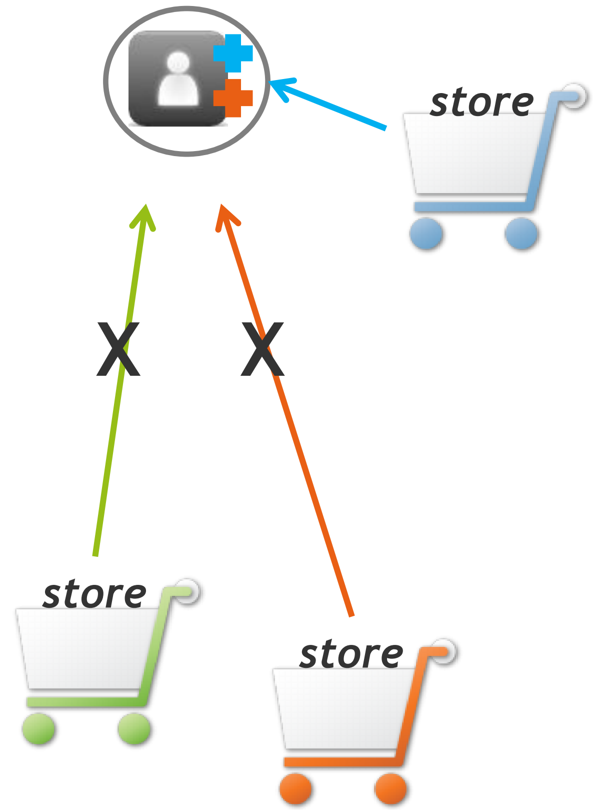

Today's preferences:

Coffee
Bookstore
Spicy

Know location and prefs

Targeted ads
Coffee close

Approach: Location-based services

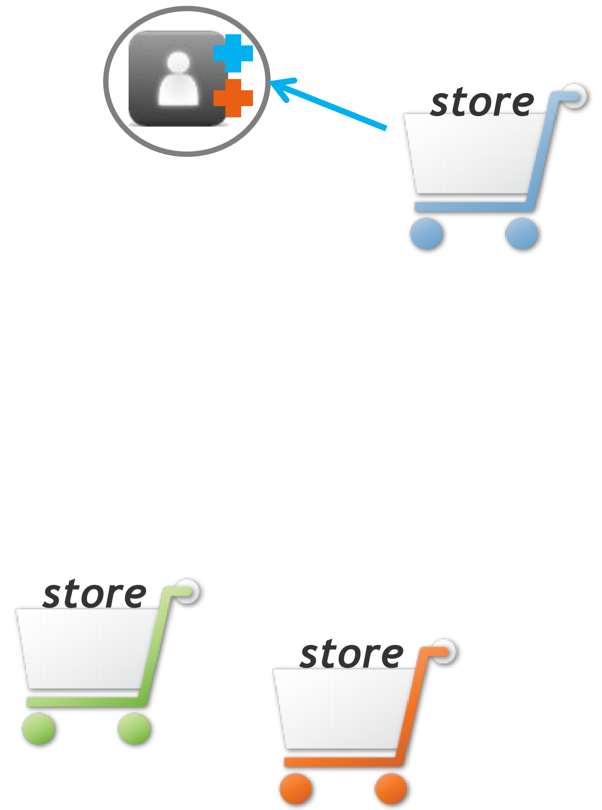

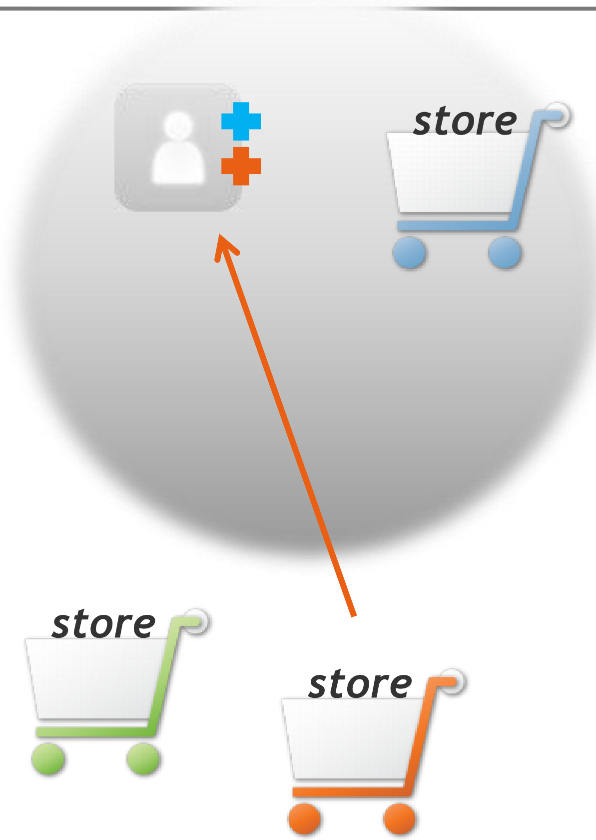

Know location and prefs

Targeted ads
Coffee close

Don't know location

Semi-targeted ads
Bookstore far

Approach: Location-based services

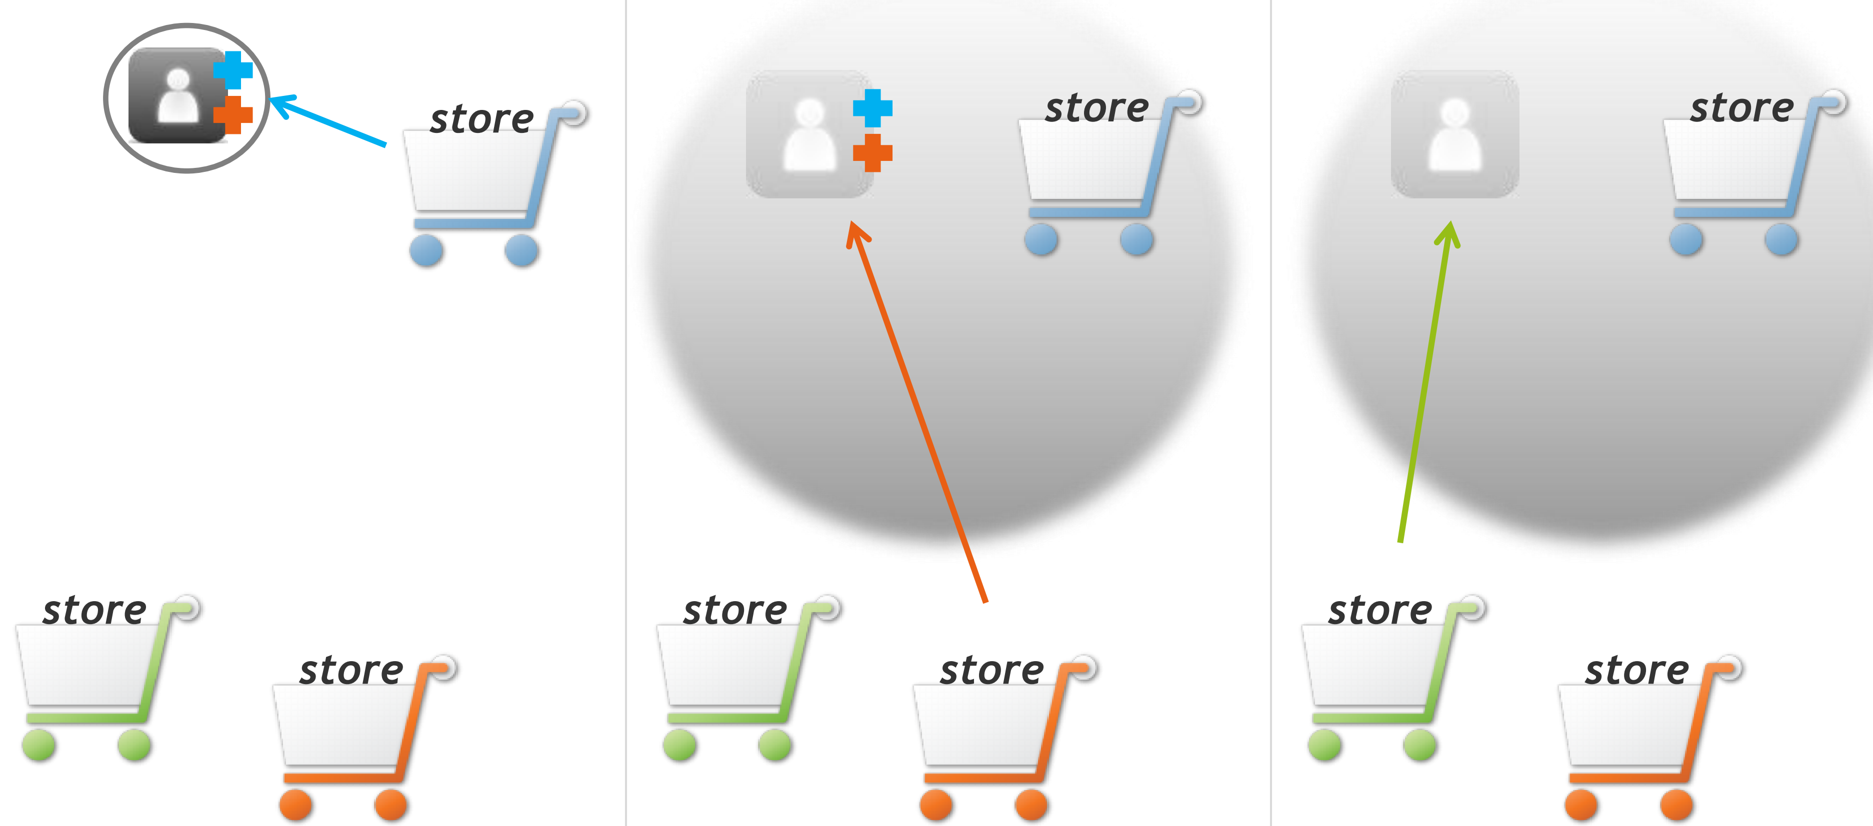

Know location and prefs

Targeted ads
Coffee close

Don’t know location

Semi-targeted ads
Bookstore far

Don’t know

Non-targeted ads
Non-spicy food far

Approach: Location-based services

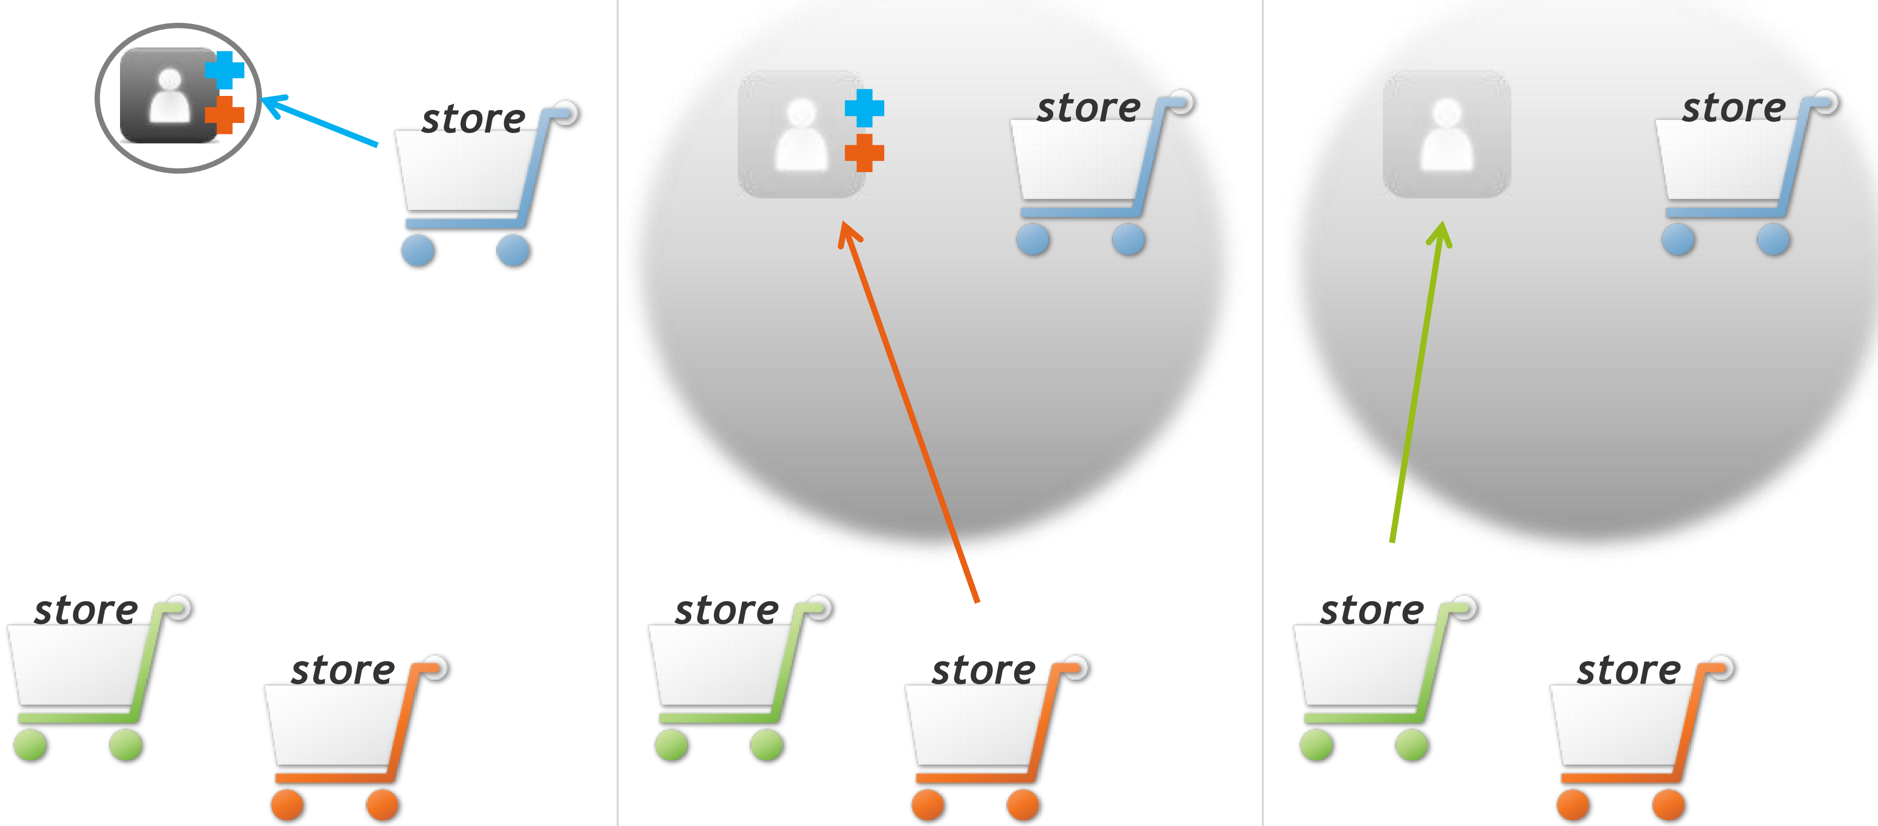

$\rho_{loc+pref}$ ρ_{pref} ρ_0

Value of location data

Value of preferences

Building the model

Complex because

- Spatially distributed users 
- Spatially distributed businesses that trigger transactions 

- Transactions depend on location and user preferences
- User location known accurately or not

Goal: new models that provide insight

- What is the value created by a knowledge of user location and/or of user preferences?
- Which one is more valuable?

Spatial processes

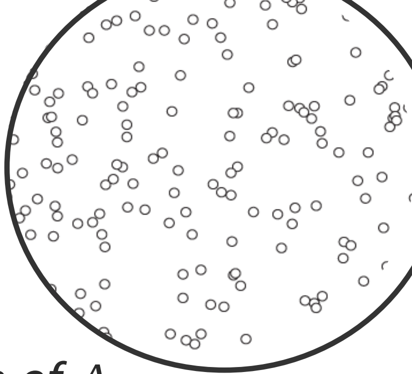

Spatial Poisson model

- Φ is a Poisson process of intensity λ on A if
  - Number of points *N(A)* is Poisson with rate λ *x surface of* A
  - Number of points in disjoint sets are independent variables

$$p(N(A)=k)=e^{-\lambda|A|}\frac{(\lambda\,|\,A|\,)^k}{k!}=e^{-\lambda\pi R^2}\frac{(\lambda\pi R^2)^k}{k!}$$

Boolean or germ-grain model

- Germs = points of Poisson process of density λ
- Grain = ball of radius R
- Prob of m-*coverage* $p(m,\lambda)=e^{-\lambda\pi R^2}\dfrac{(\lambda\pi R^2)^m}{m!}$

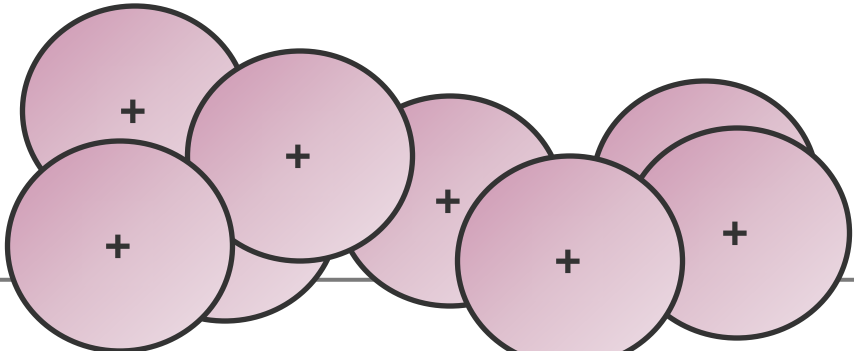

Model assumptions

Businesses

- Type n (coffee, bookstore, restaurant…) distributed according to independent spatial Poisson process λ_n . Denote $\lambda = \Sigma\lambda_k$

Users

- Spatial Poisson process of density v
- Class $(k, \mathbf{i})$ has random preference list $\boldsymbol{i}=(i_1,.. i_k)$ with prob $\pi(\mathbf{i},k)$
- Vicinity = ball of radius R

Transactions

- Users receive ads that depend on *total number of services* m in vicinity that match their list. **Propensity for users to stop** $f(m)$
- Given that user stops, *revenue or value prop to number of different services in* R - drink coffee, hang out at bookstore

Model assumptions

Businesses

- Type n (coffee, bookstore, wonton, ding
  to independent spatial Poisson pro

Users

- Spatial Poisson process of density
- Class $(k, \mathbf{i})$ has random preference
- Vicinity = ball of radius R

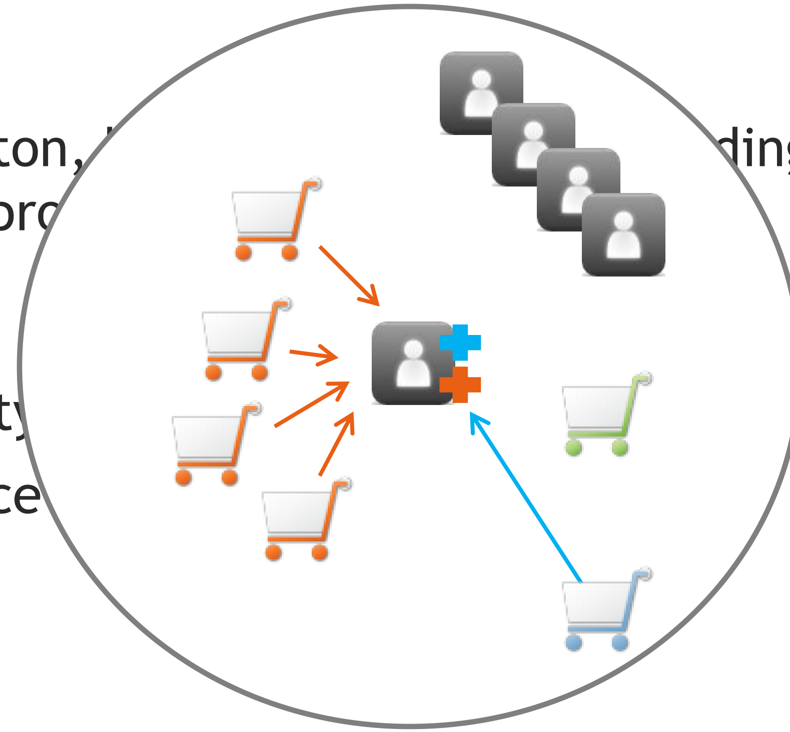

Transactions

- Users receive ads that depend on *total number of services* m in vicinity that match their list. **Propensity for users to stop** $\boldsymbol{f(m)}$
- Given that user stops, *revenue or value prop to number of different services in* R - drink coffee, hang out at bookstore

  Revenue= v x Prob (m services in R) x $f(m)$ x nb of diff services

$$\boxed{\rho} = v \sum_k \sum_i \boxed{\pi(k,\mathbf{i})} \sum_m \boxed{p(m,k,\mathbf{i})} f(m) g(m,k,\mathbf{i})$$

Case #1 – Perfect user location information

Pick a user. Given that user is of type k, $\mathbf{i}=(i_1,\dots i_k)$

- Poisson process of $\lambda(k,\mathbf{i}) = \Sigma_{j=1,..k}\lambda_{ij}$ of services present in its list
- Location m-*covered* with $p(m,k,\mathbf{i}) = e^{-\lambda(k,i)\pi R^2}\dfrac{(\lambda(k,\mathbf{i})\pi R^2)^m}{m!}$
- Mean number of different services among the m
  - No service of type p among the m $(1-\lambda_{i_p}/\lambda(k,\mathbf{i}))^m$
  - $g(m,k,\mathbf{i}) = \sum_{p=1}^{k}(1-(1-\lambda_{i_p}/\lambda(k,\mathbf{i}))^m)$

Mean revenue generated per unit space

- Location + pref $\boxed{\rho_{loc+pref}} = v\sum_k\sum_i \boxed{\pi(k,\mathbf{i})}\sum_m \boxed{p(m,k,\mathbf{i})}\, f(m)g(m,k,\mathbf{i})$
- Potential $\Phi = v\sum_k\sum_i \pi(k,\mathbf{i})\sum_m p(m,k,\mathbf{i})g(m,k,\mathbf{i})$
- Prob of stopping $p_{stop} = \sum_k\sum_i \pi(k,\mathbf{i})\sum_m p(m,k,\mathbf{i})f(m)$
- No location or pref $\boxed{\rho_0} = p_{stop}\times\Phi$

Case #2 - Imperfect user location information

User localized at distance r from true location

Case r > 2R

- Services at real location independent of services at estimated location
- Revenue with prefs, but no loc

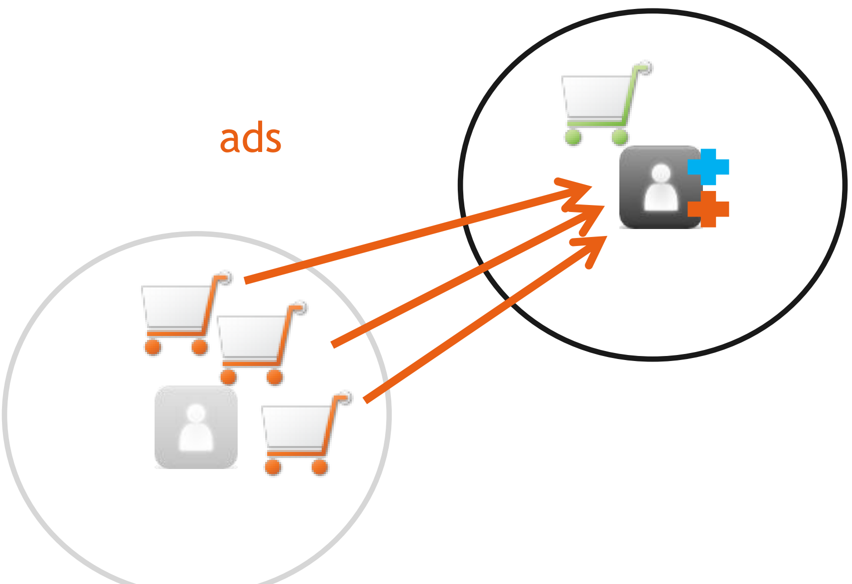

$\boxed{\rho_{pref}}$

Case r < 2R

Numerical results

Propensity to stop: $f(m) = 1 - \alpha^m$, $0< \alpha<1$

$\alpha = 0$ high propensity to react to ads or recommendations

Models psychological behavior of user

Geometric list of preferences $\pi(k,i) = \beta^k(1-\beta)\binom{N}{k}^{-1}$

$\lambda_n = \lambda$ for all n; λ is the spatial density of services

Numerical results: location vs preferences

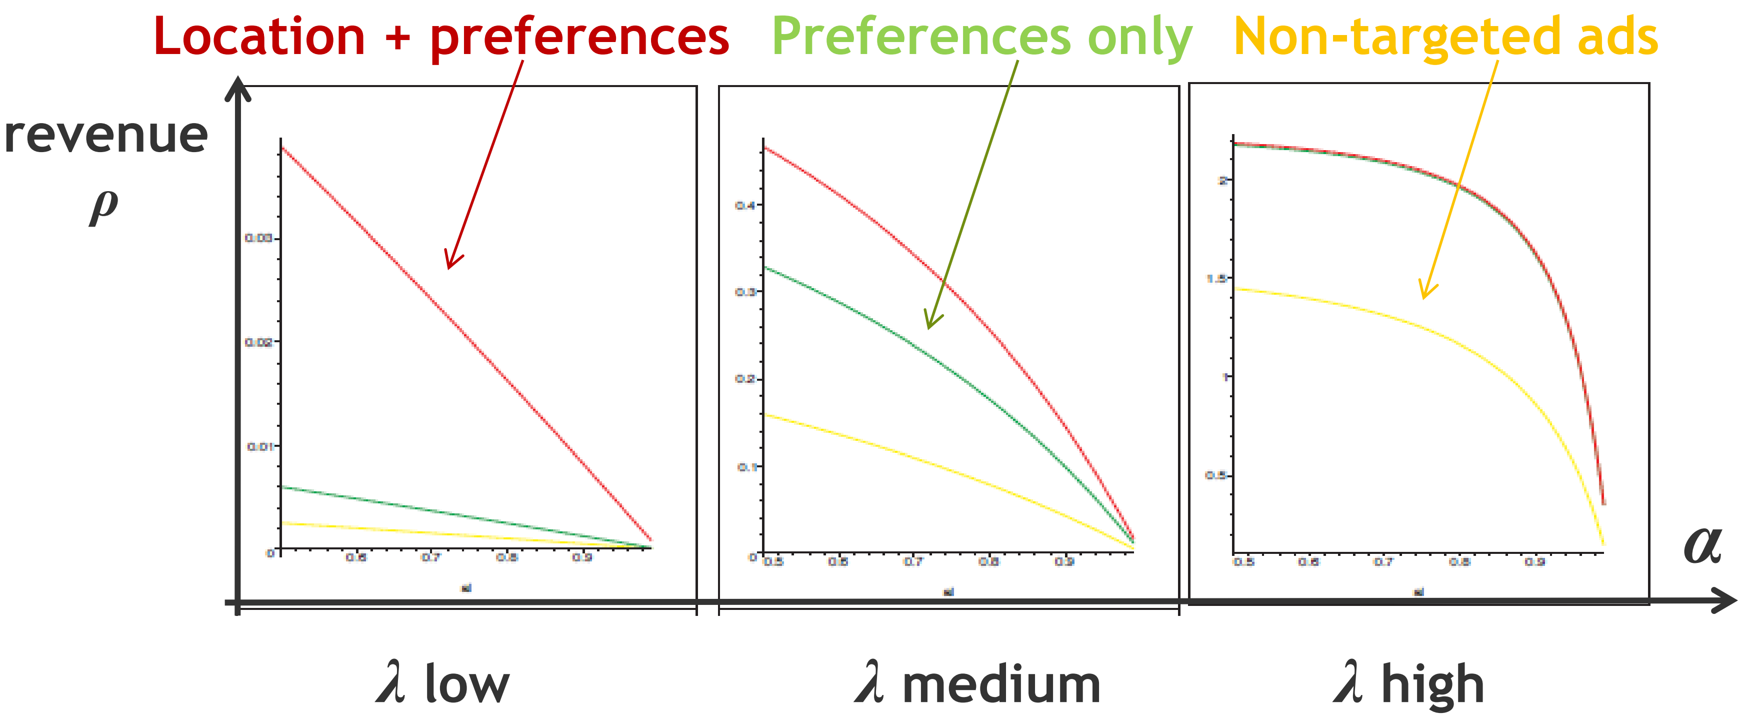

Key takeaway:

Profile data more important in dense urban cores

Location data more important in sparser areas

Simple but powerful model for location-based ads, Tinder, …

Takeaway

Stochastic geometry and stochastic processes just as important up the stack as they have been down the stack…

Will remain important given emerging trends

3. Guided usage

$f(m)$ propensity function

All interactions will be guided (Google, Amazon, yelp,..): choice, like

Rich area of research

Recommendation systems (performance, bias,..)

Impact of recommender systems on population

$$R_i(t) = sign[P_i + C_i(t) + \sum_{j \in N_j} N_{ij} R_j(t-1)]$$

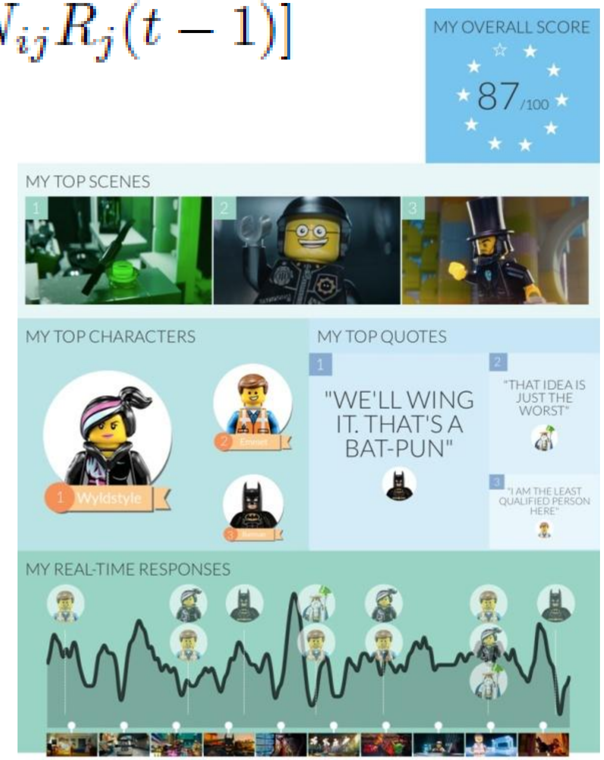

User feedaback & analysis

Impact on platform and bottom of the stack?

Itinerary - 2015

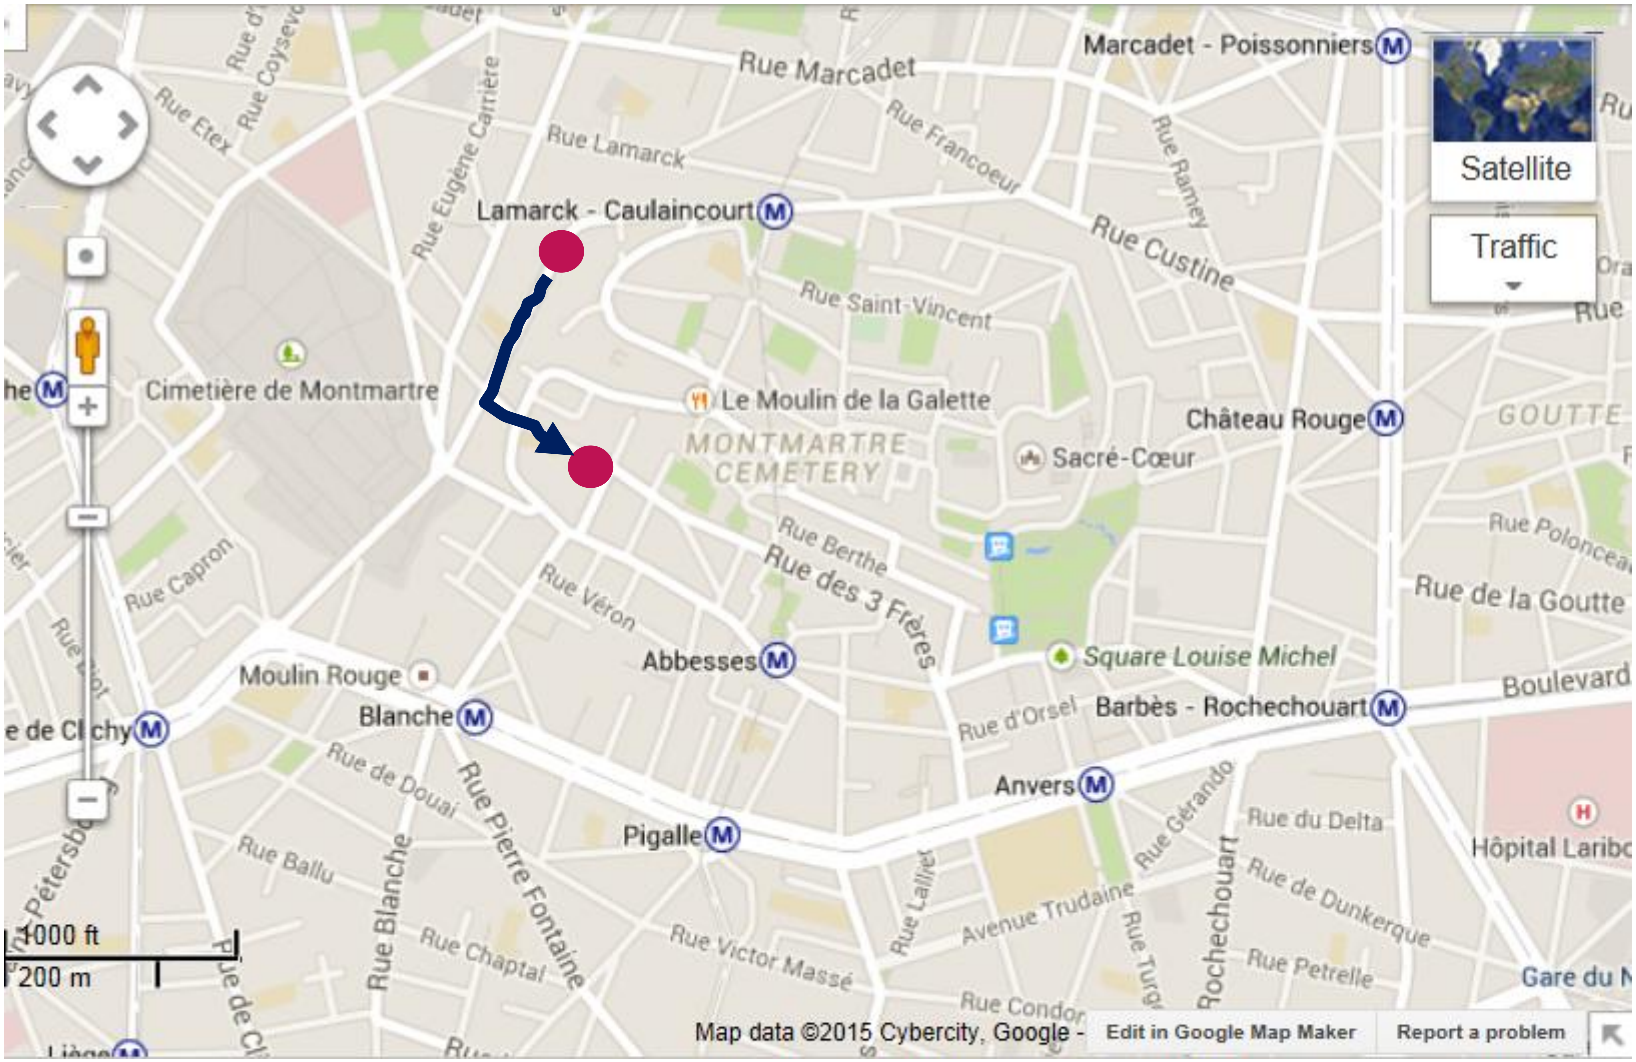

Itinerary - 2020

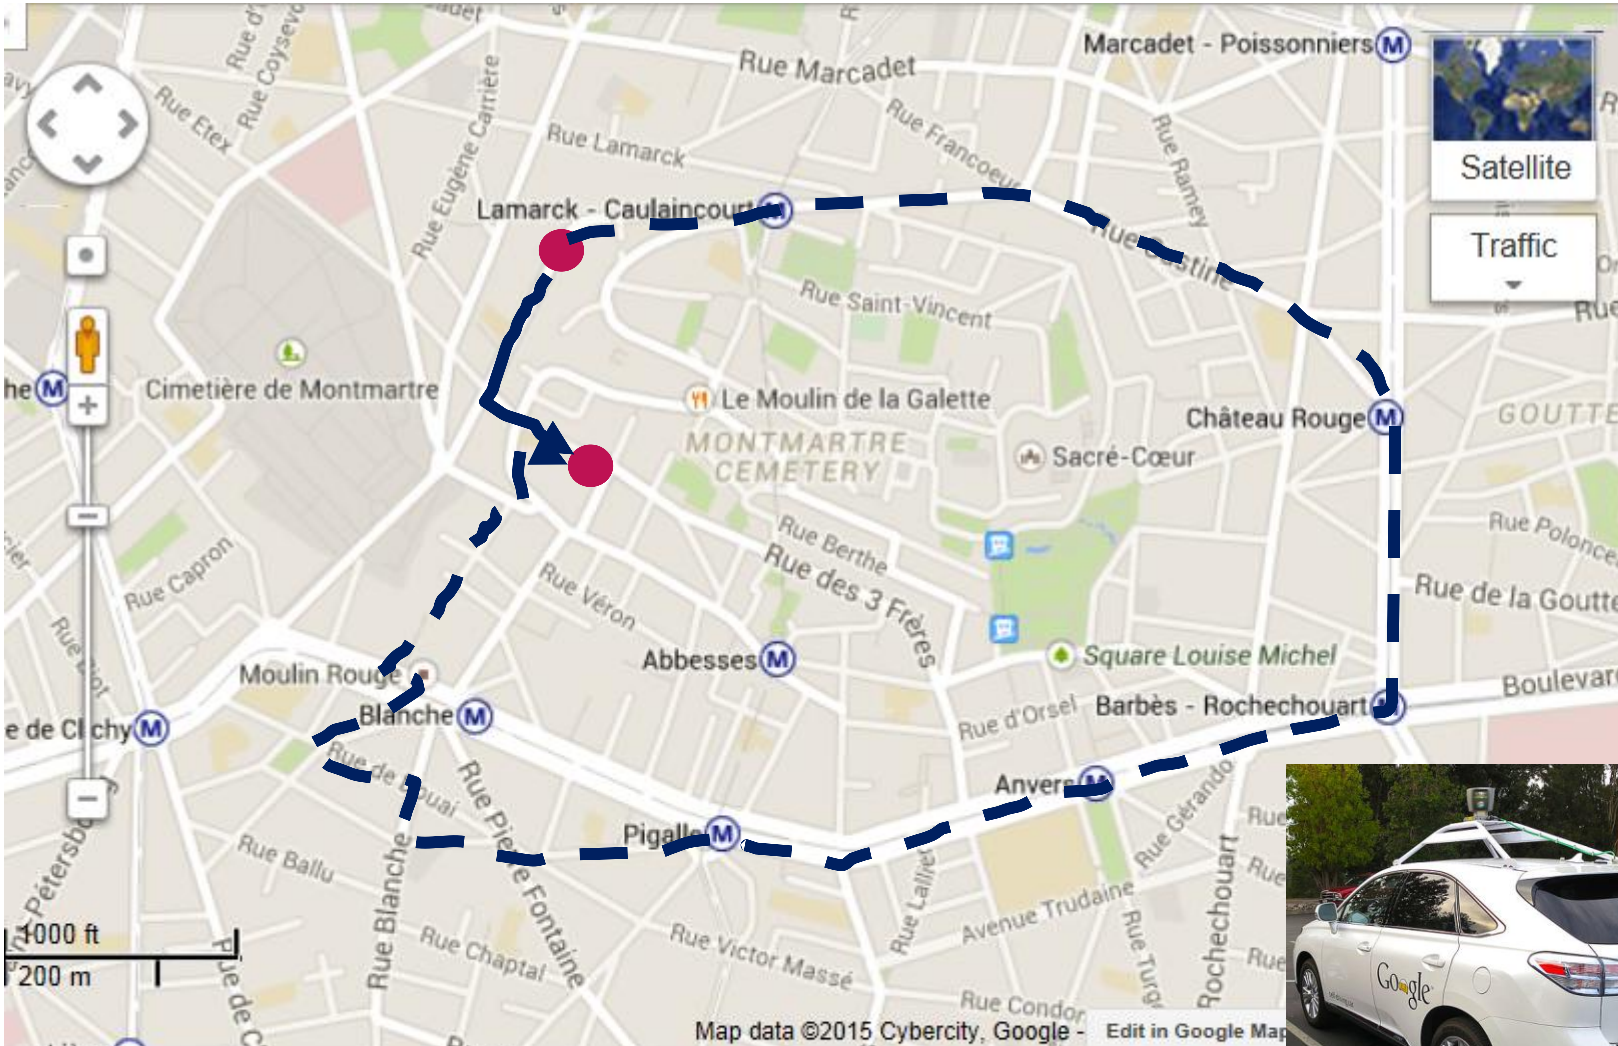

Itinerary - 2020

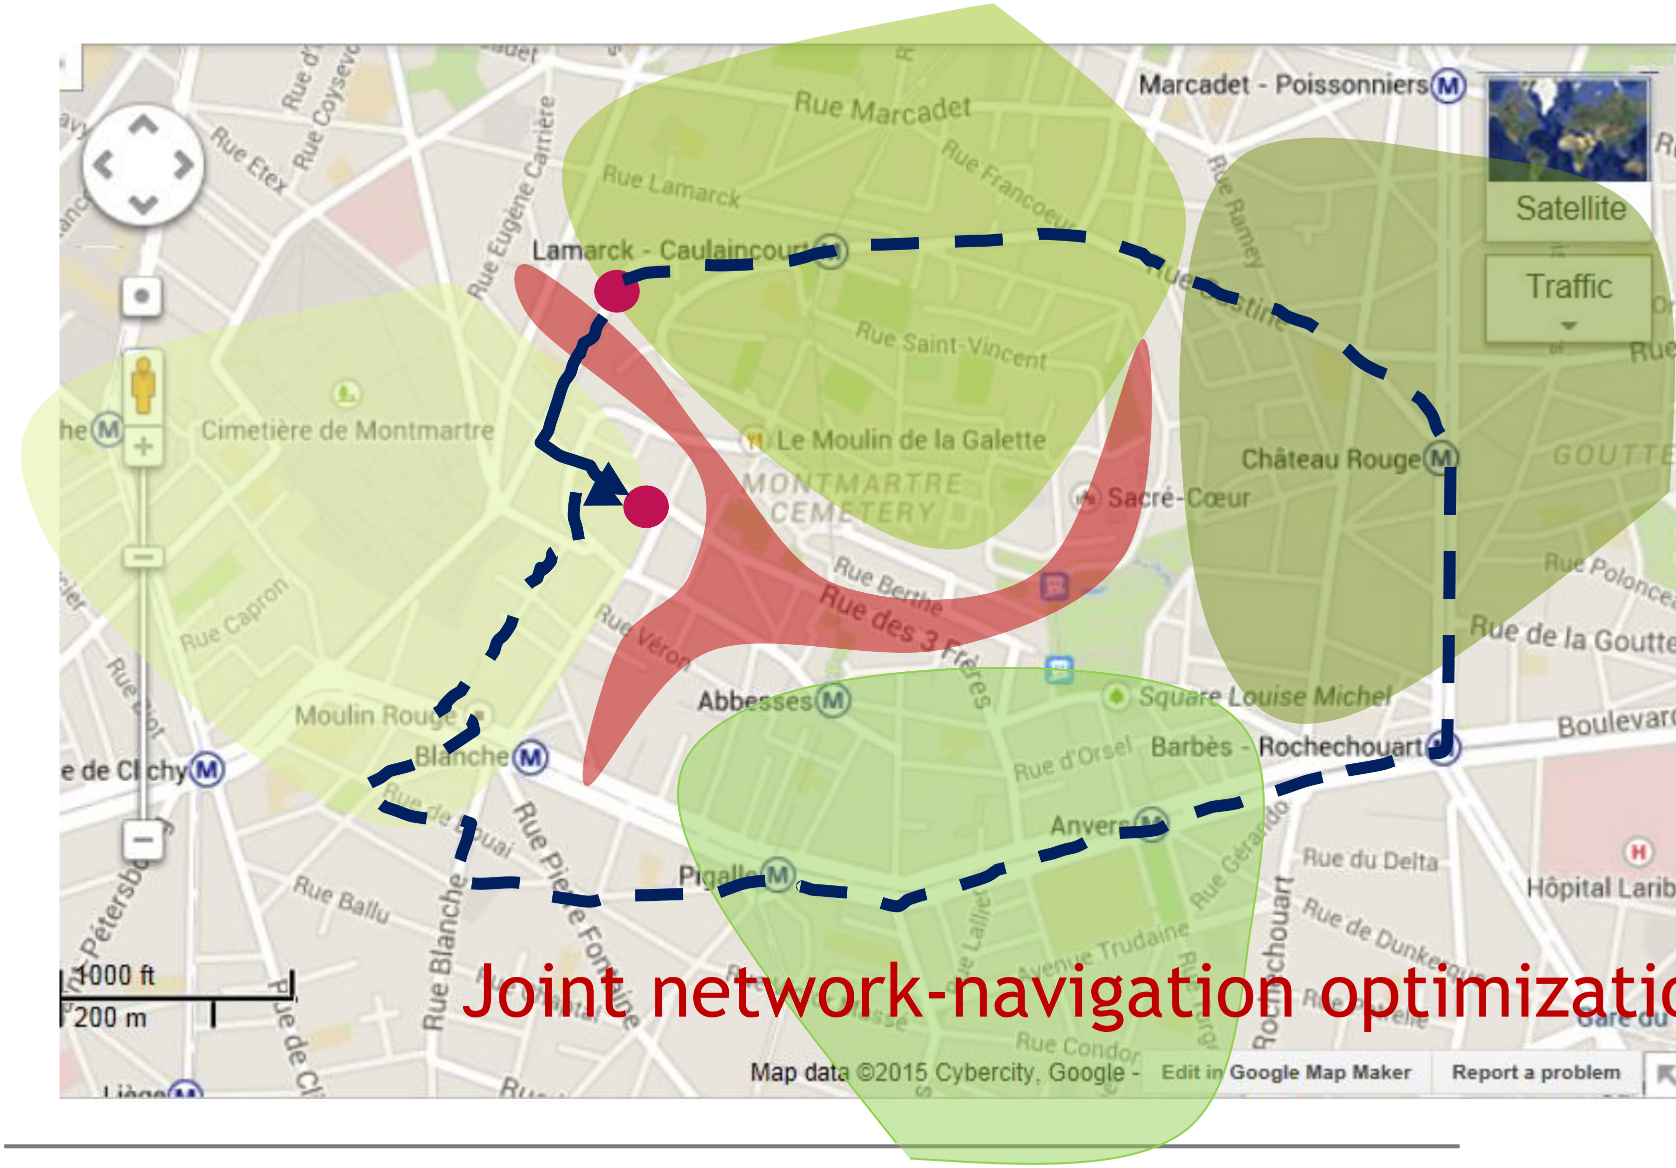

Joint network-navigation optimization

Thank you