



## INFORMS Transactions on Education

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

### Service Engineering: Data-Based Course Development and Teaching

Avishai Mandelbaum, Sergey Zeltyn,



To cite this article:

Avishai Mandelbaum, Sergey Zeltyn, (2010) Service Engineering: Data-Based Course Development and Teaching. INFORMS Transactions on Education 11(1):3-19. <https://doi.org/10.1287/ited.1100.0058>

Full terms and conditions of use: <http://pubsonline.informs.org/page/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](mailto: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.

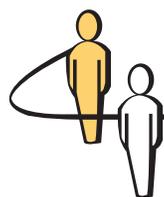
Copyright © 2010, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



# Service Engineering: Data-Based Course Development and Teaching

Avishai Mandelbaum

Faculty of Industrial Engineering and Management, Technion, 32000 Haifa, Israel, [avim@tx.technion.ac.il](mailto:avim@tx.technion.ac.il)

Sergey Zeltyn

IBM Research Lab, 31905 Haifa, Israel, [sergeyz@il.ibm.com](mailto:sergeyz@il.ibm.com)

In this exposition, we discuss empirically based teaching in the newly emerging field of “service engineering.” Specifically, we survey a service engineering course, taught at the Technion–Israel Institute of Technology. The course was “born” about 15 years ago as a graduate seminar and ultimately took its present form as a core course for the undergraduate program in industrial engineering and management. The role of measurements and data as teaching-enhancers and research-drivers is underscored. In addition, we emphasize that data granularity must reach the individual-transaction level. We describe customized databases and software tools that facilitate operational and statistical analysis of services; this includes the use of SEESat, a data-user interface that was developed at the Technion’s SEE Laboratory, for research and educational purposes. Some unique aspects of the course are the incorporation of state-of-the-art research and real-world data in lectures, recitations, and home assignments, as amply presented throughout this work. The application focus of the surveyed course has been telephone call centers, which constitute an explosively growing branch of the service industry. The course is now expanding to also cover healthcare, especially hospitals; some examples from other service areas (e.g., the justice system) are described as well.

*Key words:* teaching engineering; teaching queueing; teaching service operations management; interdisciplinary teaching; service science

*History:* Received: August 2009; accepted: July 2010.

## 1. Introduction to Service Engineering

### 1.1. Motivation and Contents of the Paper

The service sector is central in the life of postindustrial societies—more than 70% of the gross national product in most developed countries is attributable to this sector. In concert with this state of affairs, there exists a growing demand for high-quality multidisciplinary research in the field of services, as well as for a significant number of *service engineers*, namely, scientifically educated specialists that are capable of designing service systems, as well as solving multifaceted problems that arise in their practice. (Readers are referred to [Fitzsimmons and Fitzsimmons 2004](#), especially Part 1, for background on services, including some distinguishing characteristics of services relative to manufacturing.)

In the United States, the academic home for the area of services has traditionally been the business school, where services have been taught most often as “service marketing.” Another business school option

is “service management,” either within operations management courses or, rather rarely, as a standalone service operations course. As an engineering discipline, the natural home for services are industrial engineering units. Indeed, our original use of the term “service engineering” was conceived by combining the relevant words in “service management” and “industrial engineering.” We had roughly in mind a “*New Age Industrial Engineer* that must combine technological knowledge with process design to create the delivery systems of the future ([Frei et al. 1998](#), p. 33).”

It is our belief that there exists a broad gap between academia’s supply and the demand for service science and engineering. Focusing on the education, universities either do not offer service engineering courses or, when they do so, the education quality is typically not at par with that of the traditional engineering disciplines, in particular that provided by leading industrial engineering units.

The goal of the *service engineering* (ServEng) course, “born” at the Technion–Israel Institute of Technology

about 15 years ago (Service Engineering 2010), is to narrow down the educational gap between demand and supply while also stimulating service research. As we perceive it, the ultimate goal of service engineering is to *develop scientifically based design principles and tools (often culminating in software) that support and balance service quality, efficiency, and profitability from the likely conflicting perspectives of customers, servers, managers, and often, society also*. The goal of our ServEng course is more restricted: we take an operations point of view, hence we focus on *operational service quality*, on *service efficiency*, and on the tradeoffs between the two.

Our ServEng course takes a *data-based approach to teaching*. Thus, throughout the paper, course contents will be illustrated with the help of practical examples, drawing from data-based case studies. Students encounter such case studies continuously during lectures, recitations, and homework assignments. For example, a concept can be illustrated in a recitation via data from bank tellers' service, homework practices the concept on call center data, and students are tested on it using data from an emergency department.

The remaining part of §1 is dedicated to our paradigm of the service engineering discipline. Specifically, §1.2 presents the course homepage, where the ServEng materials can be downloaded. Section 1.3 introduces service networks as our modeling framework for a service operation, with queueing theory and science constituting the main theoretical foundation. Section 1.4 elaborates on relevant milestones in the evolution of queueing theory. Then, §1.5 focuses on the main challenges, goals, and methods in the ServEng discipline, in the context of the course. In this part, we also explain the need for approximations (compromises) in modeling and analysis of service systems.

Section 2 elaborates on the goals of the ServEng course and provides a brief survey of its history. Section 3 explains our data-based teaching approach and introduces the DataMOCCA (Data Model for Call Center Analysis) system: through its graphical user interface, SEESat, this system enables a friendly effective access to several large data repositories from service systems (currently, call centers and emergency departments), which have been used for research and teaching (specifically in lectures, recitations, and homework of the ServEng course).

In §4 we briefly describe the lectures of the course. Section 5 lists the data-based homework assignments, where students must use DataMOCCA/SEESat tools for the analysis of call center data. We conclude, in §6, with acknowledgments to those who helped bring the course to where it is now, intertwining it with some additional course history.<sup>1</sup>

<sup>1</sup> A "full version" of this paper can be downloaded from the journal website. The full version includes expanded descriptions of lectures

## 1.2. Course Homepage(s)

Before continuing, readers are advised to browse through the ServEng website and get an idea of its structure and contents. References to relevant links will be provided as we go along. For example, the "Table of Contents" of the recently taught course is in <http://ie.technion.ac.il/serveng/lectures/schedule1.doc>.

The formal course website is <http://ie.technion.ac.il/serveng>, which includes the material of a *complete* course—the last one to have been taught. Then, as the course is being offered, a second website is constructed gradually and updated as the course progresses; the address of this dynamic website is appended by the course's year and semester. For example, the site of the course taught during the spring semester of 2010 has the address <http://ie.technion.ac.il/serveng2010S>. (The earlier winter semester had 2010W appending its address.)

## 1.3. Service Networks: Models of Congestion-Prone Service Operations

The title of this section reflects our angle on service operations—we often view them as stochastic (random) or deterministic (fluid) systems, within the operations research paradigm of queueing networks. To support this view, let us first present our conception of the roles of queues in services from the perspectives of customers, servers, and managers. We shall then describe service networks, continuing with some relevant queueing theory history.

**1.3.1. On Queues in Service.** *Queues in services* are often the arena where customers, service-providers (servers), and managers interact (establish contact) to jointly create the service experience. Processwise, queues play in services much the same role as inventories in manufacturing (see JIT = just-in-time, TBC = time-based competition, etc.). However, in addition, "human queues" express preferences, complain, abandon, and even spread around negative impressions. Thus:

- *Customers* treat the queueing experience as a window to the service-providing party through which their judgment of it is shaped for better or worse.
- *Servers* can use the queue as a clearly visible proxy for the state of the system, based on which, among other things, service protocols can be exercised (e.g., customers' priorities).
- *Managers* can use queues as indicators (queues are the means, not the goals) for control and improvement opportunities. Indeed, queues provide unbiased quantifiable measures (these are not abundant in services) in terms of which performance is relatively easy to

and some homework assignments, as well as an additional section, "The Fusion of Research and Teaching."

monitor and goals (mainly tactical and operational but sometimes also strategic) are naturally formulated.

Our point of view is thus clear: the design, analysis, and management of queues in service operations could and should constitute a central driver and enabler in the continuous pursuit of service quality, efficiency, and profitability.

### 1.3.2. On Service Networks and Their Analysis.

*Service networks* here refer to *dynamic (process) models* (mostly analytical, sometimes empirical, and, rarely, simulation) of a service operation as a *queueing network*. The dynamics is that of serving *human customers*, either directly face-to-face or through phone calls, email, the Internet, etc. Informally, a *queueing network* can be thought of as consisting of interconnected service stations. Each station is occupied by servers who are dedicated to serve customers queued at the station.

In its simplest version, the evolution of a service station over time is stationary, as statistically identical customers arrive at the station either exogenously or from other stations. Upon arrival, a customer is matched with an idle server, if there is any; otherwise, the customer joins a queue and gets served on a first-come, first-served (FCFS) basis. Upon service completion, customers either leave the network or move on to another station in anticipation of additional service. Extensions to this simplest version cover, for example, models with nonstationary arrivals (peak-loads); multitype customers that adhere to alternative service and routing protocols; customers' abandonment while waiting (after losing their patience); finite waiting capacities that give rise to blocking, splitting (fork), and matching (join) of customers; and much more.

In analyzing a service network, we find it useful to be guided by the following four steps (although, unfortunately, most often, only the first three are applied or even applicable):

- *Can we do it?* Deterministic capacity analysis, via process-flow diagrams (spreadsheets, linear programming), which identifies resource-bottlenecks (or at least candidates for such) and yields utilization profiles.
- *How long will it take?* Typically, stochastic response-time analysis, via analytical queueing network models (exact or approximate) or simulations, which yields congestion curves. Note that when predictable variability prevails and dominates, then the "fluid view" and the corresponding deterministic methods are appropriate; see §4.4 of Full Version (Mandelbaum and Zeltyn 2009).
- *Can we do better?* Sensitivity and parametric ("what-if") analysis of measures of performance (MOPs) or scenarios, which yields directions and magnitudes for improvements.

- *How much better can we do?* or, put simply, "What is optimal to do?" This type of questions has been typically addressed via optimal control (exact or asymptotic) which, as a rule, is difficult but becoming more and more feasible. (Another developing research direction is optimization via simulation—see Feldman 2008 for a call-centers example, which is now in the process of being incorporated into ServEng.)

Several examples of service networks are covered in Full Version (Mandelbaum and Zeltyn 2009); see, for example, the dynamic-stochastic (DS) PERT/CMP networks (sometimes referred to as fork-join or split-match networks) in §4.3.

### 1.4. Some Relevant History of Queueing Theory

Queueing theory provides a central mathematical foundation for ServEng. Accordingly, a variety of queueing models, including Markovian, fluid, and heavy-traffic, are studied in our course. It is thus appropriate to cover some relevant milestones in this theory's evolution.

**1.4.1. The Early Days.** The father of queueing theory was the Danish telecommunication engineer Agner Krarup Erlang who, around 1910–1920, introduced and analyzed the first mathematical queueing models. Erlang's models (Erlang 1948) are standardly taught in elementary/introductory academic courses (for example,  $M/M/n$ ,  $M/M/n/n$ ), as they are still cornerstones of today's telecommunication models (where  $M/M/n/n$  is known as Erlang-B, "B" apparently for blocking—the central feature of this model—and  $M/M/n$  is referred to as Erlang-C, "C" conceivably because it is subsequent to "B"). Moreover, and more relevant to our present discussion,  $M/M/n$  is still the workhorse that supports workforce decisions in many telephone call centers; see §4.6 of Full Version (Mandelbaum and Zeltyn 2009).

Another seminal contributor to queueing theory, Scandinavian (Swedish) as well, is Conny Palm, who in 1940–1950 added to Erlang's  $M/M/n$  queue the option of customers' abandonment (Palm 1957). We shall refer to Palm's model as Palm/Erlang-A, or just Erlang-A for short (unfortunately, and rather unjustly, to Palm, but Erlang was "there first"). The "A" stands for abandonment and for the fact that Erlang-A is a mathematical interpolation between Erlang-B and Erlang-C. Palm, however, has been mostly known for his analysis of time-varying systems. This analysis is also of great relevance to service operations and, hence, was covered, in some way, in ServEng.

A next seminal step (one might say a "discontinuity" in the evolution of queueing research) is due to James R. Jackson, who was responsible for the mathematical extension of Erlang's model of a *single* queueing station to a system of *networked* queueing stations, or queueing networks, around 1955–1965.

Jackson was motivated by manufacturing systems and actually analyzed open and semiopen networks. Closed networks, relevant to healthcare and call centers with an answering machine, as it turns out, were analyzed in the mid-1960s by William J. Gordon and Gordon F. Newell. Interestingly, Newell was a transportation engineer at Berkeley and was also the earliest influential advocate of incorporating fluid models as a standard part of queueing theory; see his textbook (Newell 1982). A student of Newell, Randolph W. Hall, who is currently a Professor at the University of Southern California, wrote an excellent queueing book (Hall 1991) that has influenced our teaching of service engineering; Hall is currently working on healthcare systems, adopting the fluid-view (described below) to model the flows of patients in hospitals; see Hall (2006).

Jackson networks are the simplest theoretically tractable models of queueing networks. (Their simplicity stems from the fact that, in steady state and at a fixed time, each station in the network behaves like a naturally corresponding birth-death model, *independently* of the other stations.) The next step beyond Jackson networks are BCMP/Whittle/Kelly networks, where the heterogeneity of customers is acknowledged by segregating them into classes. However, service operations often exhibit features that are not captured by Jackson and BCMP/Whittle/Kelly networks. Further generalizations are therefore needed, which include precedence constraints (fork-join, or split-match networks), models with one-to-many correspondence between customer types and resources (skills-based routing, agile workforce), and models that exhibit transient behavior (as opposed to steady-state analysis).

**1.4.2. QED Queues.** The key tradeoff in running a service operations is that between service efficiency and quality, which queueing models are ideal to capture. This tradeoff is most delicate in large systems (many servers), but here exact analysis of queueing models turns out to be limited in its insight. This was already recognized by Erlang around 1910, and later by Pollaczek around 1930, both of whom resorted to approximations. However, the first to put operational insights of many-server queues (as relevant to us) on a sound mathematical footing were Shlomo Halfin and Ward Whitt in the early 1980s, in the context of Erlang-C (and its extension GI/M/n). They introduced what we shall call QED queues (Halfin and Whitt 1981), which stands for queues that are *both* quality- and efficiency-driven, hence, their name.

QED queues emerge within an asymptotic framework that theoretically and insightfully supports the analysis of the efficiency-quality tradeoff of *many-server* queueing systems. The theory culminates in

a simple rule of thumb for staffing, the square-root staffing rule: If the offered load is  $R$  Erlangs ( $R =$  arrival rate times average service time), then

$$n \approx R + \beta\sqrt{R}$$

is a staffing level that would appropriately balance quality and efficiency. Here  $\beta$  is a constant, which corresponds to grade-of-service; its specific value (which is relatively small—less than 1.5 in absolute value) can be determined by economic considerations (for example, the ratio between delay costs and staffing costs).

Prime examples of QED queues are well-run telephone call centers, but to properly model these, one must generalize the Halfin-Whitt framework to allow for customers' impatience. This was done in a Technion M.Sc. thesis by Ofer Garnett, in the late 1990s and was later published in Garnett et al. (2002). At that same time, a generalization of the Halfin-Whitt framework to time-varying Jackson-like networks was carried out with Bill Massey and Marty Reiman (Mandelbaum et al. 1998, under the name Markovian Service Networks). In analogy to QED queues, there are also ED (efficiency-driven) and QD (quality-driven) queues, both arising from asymptotic analysis as well. Erlang-C, in these three operational regimes, was treated in Borst et al. (2004); generalizations to Erlang-A and relatives is the subject of the Ph.D. thesis of the second author (S.Z), published in Zeltyn and Mandelbaum (2005). The research part of the website of Ward Whitt is recommended for further references on QED/ED/QD queues (Whitt 2010).

QED approximations turn out to be amazingly accurate and robust. This arose explicitly in Borst et al. (2004), where *square-root staffing* was shown to be an asymptotically optimal staffing level ( $n$  in an  $M/M/n$  queue). It turns out that the square-root recipe rarely deviates from the actual optimal value by more than one server, over a wide range of staffing levels—from the very few (10 or so) to the very many (1000s). This unexpected accuracy has been recurring for many other models since, and it can now be explained theoretically for the Erlang C/B/A models, as well as  $M/D/n$ ; see, for example, Janssen et al. (2010). The accuracy of QED approximations has practical significance because it expands the scope of these approximations to relatively small systems—in particular health-care systems (e.g., emergency departments and internal wards); this was first observed by Jennings and de Vericourt (2007). It is further expanded on in the ongoing Ph.D. thesis of Yom-Tov (2007), which supports decisions on static capacity (hospital beds) and dynamic capacity (e.g., nurses). Interestingly, Galit's basic queueing model for beds + nurses is the one that also appeared in the M.Sc. thesis of Khudyakov (2006) on call centers with an answering machine—indeed, beds correspond to trunk lines and nurses to telephone agents.

**1.4.3. ED Queues.** A balance between service quality and efficiency is often desirable, yet sometimes it is infeasible, or perhaps not even optimal. Then one settles for an ED (efficiency-driven) service system, in which the focus is on high utilization of resources (the servers), at the cost of service quality (relatively long delays). In call centers, ED performance arises when resources are restricted (e.g., government services); however, with large enough of a system, service level could still be acceptable or even better (see Ward Whitt’s homepage (Whitt 2010) for theoretical papers on many-server ED queues). However, ED performance is in fact prevalent in healthcare, which renders small-server heavy-traffic asymptotics more relevant—see, for example, the ongoing M.Sc. thesis of Zviran (2008) on fork-join queues in heavy traffic.

**1.4.4. Summary.** To summarize, queueing networks have been successfully used to model systems of manufacturing, transportation, computers, and telecommunication. For us they serve as indispensable models of service systems, in which customers are human and queues, broadly interpreted, capture prevalent delays in the service process. The service interface could be phone-to-phone (naturally measured in units of seconds) or face-to-face (in minutes), fax-to-fax (hours), letter-to-letter (days), face-to-machine (e.g., ATM, perhaps also Internet), etc. The finer the time-scale, the greater is the challenge of design and management. Accordingly, the greater is the need for supporting rigorous models, a need that further increases with scale, scope, and complexity.

**1.5. Service Engineering: Challenges, Goals, and Methods**

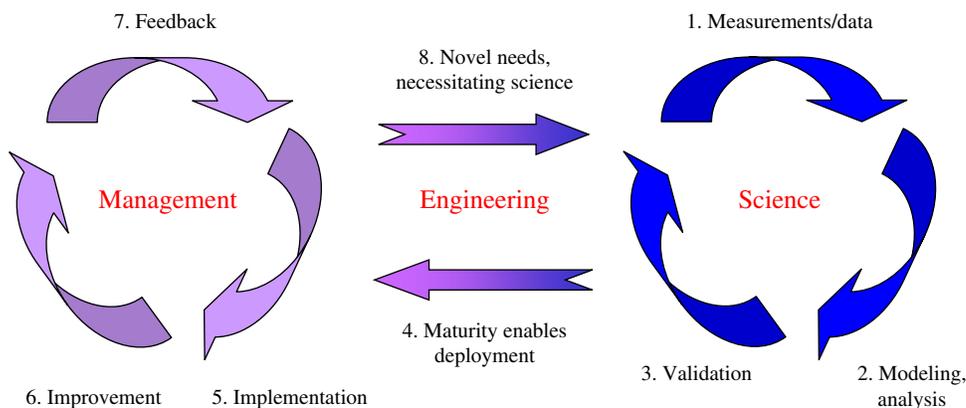
As already mentioned, we have been advocating the terminology “service engineering” to describe our research, teaching, and consulting on service systems. Based on our experience, we now elaborate on the role of the service engineer as a natural mediator between the service scientist and the service manager.

**1.5.1. Our Service Science and Service Engineering Paradigm.** Figure 1 describes our view on the role that service engineering plays relative to the classical scientific cycle. This cycle starts with measurements that gather data, which is used to create models: mathematical, simulation, or sometimes merely conceptual ones. The models are then validated, which might lead to further measurements and model refinements, continuing within this cycle until a maturity level is reached that enables deployment. Now service engineering transforms this science to applications, via design. Then management takes over, implementing the new designs which, hopefully, leads to improvements. Feedback from the application often generates novel needs, which might necessitate further measurements, further models, and so on. This scientific paradigm of measurement, modeling, and validation (which is extended here to engineering and management), is axiomatic in natural sciences (biology, physics), but it is definitely a newcomer to the “science of the artificial,” as Simon (1996) called it—that is, the science of man-made systems, such as service systems. Being mathematicians in first training, we were certainly not brought up with this thinking, but service science naturally calls for this approach, and we are thus learning.

**1.5.2. Challenges and Goals.** Research, teaching and practice of service engineering, as we perceive it, should take a *designer’s view*. Design challenges pertain, for example, to the following issues.

- Service strategy: determinants of service-quality levels, full-service versus self-service, customization versus standardization, warranty (after-sales support depth), etc.
- Choosing the service interface (channel): assume that service can be potentially performed by phone and/or by email, fax, letter, or perhaps face-to-face.
- Designing service process: should one use front-office or back-office (or possibly both), should the tasks be performed sequentially or in parallel, etc.

Figure 1 Service Science Paradigm



Downloaded from informs.org by [1.198.223.170] on 12 February 2018, at 06:00 . For personal use only, all rights reserved.

Additional information, including supplemental material and rights and permission policies, is available at <http://ite.pubs.informs.org>.

- Control issues: which customers to admit to service, priority scheduling, skills-based routing, exploiting idleness, etc.
- Operational resource management: how many servers to staff to reach an acceptable service level (with staffing being performed offline or online, the latter as a corrective action to inaccurate forecasting), shift scheduling, etc.
- Designing the service environment: waiting experience, using busy signal versus music in call centers, providing information to customers (e.g., predicting delay durations), etc.
- Marketing: customer segmentation, cross- or up-selling, marketing-operations interfaces, etc.
- Information systems: data-base design of call-by-call operational and business data, offline and online queries, etc.
- Taking into account human factors: career paths, incentives, hiring policies, number of call center agents who are physically present at the workplace versus actual workforce level (number of agents who are actually available to take calls).

The above designer's view supports our vision of the ultimate goal of service engineering (both in research and teaching), as stated in §1.1: to develop scientifically based design principles and tools that support and balance service quality, efficiency, and profitability. We find that queueing network models constitute a natural convenient nurturing ground for the development of such principles and tools. However, the existing supporting (queueing) theory has been somewhat lacking, as will now be explained.

**1.5.3. Scientific Perspective.** The bulk of what is called queueing theory (recall §1.4) consists of research papers that formulate and analyze queueing models with realistic flavor. Most papers are knowledge-driven, where “solutions in search of a problem” are developed. Other papers are problem-driven, but most do not go far enough to a practical solution. Only relatively few articles develop theory that is either rooted in or actually settles a real-world problem, and very few carry the work as far as validating the model or the solution. In concert with this state of affairs, not much is available of what could be called “*queueing science*,” or perhaps the science of congestion, which should supplement traditional queueing theory with data-based models, observations, and experiments. In service networks, such “science” is lagging behind that in telecommunications, transportation, computers, and manufacturing. Key reasons seem to be the difficulty to measure services (any scientific endeavor ought to start with measurements), combined with the need to incorporate human factors (which are notoriously difficult to quantify). Because reliable measurements ought to constitute a prerequisite for proper management (see

TQM = total-quality management, for example), the subject of measurements and proper statistical inference is important in our context.

**1.5.4. Engineering Perspective.** Service networks provide a platform for advancing what could be described as queueing science and management engineering of sociotechnical systems. Management engineering links management science with management practice by “solving problems with existing tools in novel ways” (Corbett and Van Wassenhove 1993). Quoting the late R. Herman (1992), acknowledged as the “father of transportation science,” noted that sociotechnical systems are to be distinguished from, say, physical and engineering systems, as they can exhibit incredible model complexity due to human beings expressing their microgoals. (Significantly, Herman's models of complexity were nevertheless tractable through remarkable collective effects; in other words laws of large numbers which, for services as well, turn out to play a central explanatory role.) The approach and terminology that we have been using, namely, *service engineering*, is highly consistent with the once influential BPR (business-process-reengineering) evolution, as well as with ERP (enterprise-resource-planning) and CRM (customer-relations-management, or the more rational acronym *customer-revenue-management*), placing heavy emphasis on the process view and relying heavily on the accessibility of information technology.

**1.5.5. Phenomenology, or Why Approximate.** Service systems often operate over finite-time horizons (the notion of steady state then requires reinterpretation). They use heterogeneous servers, whose service capacities are time and state-dependent. Their customers are “intelligent” and typically but not always prefer short queues; they jockey, renege and, in general, react to state changes and learn with experience. Finally, service systems suffer from high variability—both predictable and unpredictable, as well as diseconomies of scale—when being decentralized and inefficient (e.g., often FCFS is the only option). Such features render the modeling of service networks a challenge and their exact analysis a rarity. This leads to research on *approximations*, typically short- but also long-run fluid and diffusion approximations. Approximations also enhance exact analysis by simplifying calculations and exposing operational regimes that arise asymptotically.

The “ultimate products” of approximations are scientifically based, practically useful rules of thumb. Examples of such rules of thumb will be demonstrated below. For example, in §4.7 of Full Version (Mandelbaum and Zeltyn 2009) on staffing, we cover

staffing rules, customized to the QD/ED/QED operational regimes; in particular, square-root staffing corresponds to the QED regime. (Recall the discussion of these regimes in §1.4.)

Our approximations also reveal the “right” scale (time and space) at which phenomena evolve. For example, in ED many-server systems, waiting time is expected to be on the order of a service time. In QED many-server systems, waiting time is expected to be one order of magnitude less than service time. To be concrete, in QED call centers, waiting is measured in seconds versus service times in minutes; in transportation, time to seek parking in a busy city downtown is measured in minutes while parking time is measured in hours; and in emergency departments, waiting to be hospitalized is measured in hours versus hospitalization time in days. As another scaling example, QED call centers enjoy abandonment rates that are inversely proportional to the square root of the number of agents; this enables one to draw performance curves, as in Figure 27 in Full Version, where a *single* graph predicts abandonment rates for call centers of *all* sizes.

## 2. Course Goals, History, and Prerequisites

The service engineering (ServEng) course has been taught for nearly 15 years at the Faculty of Industrial Engineering and Management (IE&M) in the Technion–Israel Institute of Technology. It started as a seminar for graduate students, entitled “Service Networks.” It had then evolved through the stages of a graduate course that can be attended by advanced undergraduate students to an elective undergraduate course attended also by graduate students who seek an introduction to the field of service engineering. Currently, service engineering is an undergraduate *core* course, taught each semester and attended by approximately 100 students yearly. In addition, many of the IE&M graduate students in stochastic operations research and statistics take it as well. (In general, gradual transition from an elective to a core course seems to us appropriate when a new course in the field is designed.) The ServEng website ([Service Engineering 2010](#)) is designed to contain all course materials (lecture notes/slides, recitations, homework). This website is supplemented by related research papers, slides of seminars, software, databases, and more.

Prerequisites for our course include core undergraduate courses in calculus, probability theory, statistics, stochastic processes, and optimization (linear programming). Students of our department usually take all of these courses during their first two years of studies.

As discussed in §1.1, the service sector is dominating the economics of developed countries. Still, prior

to the launch of the ServEng course, Technion IE&M students had been exposed mainly to methods and techniques inspired by manufacturing applications. (This situation seems to prevail also among IE departments at other universities.) The ServEng course aims at filling this gap by providing students with appropriate models and tools for design, operation, and analysis of service systems. Our teaching approach is data-oriented: examples from various service sectors are presented at lectures, recitations, and homework assignments, with the call center industry being the central (but in no way the only) application area.

In 15 years of course development, our service system data repositories have been developed in parallel with the theoretical material of the course. In the next section, we describe the service system data that supports the course.

## 3. Data—A Prerequisite for Research and Teaching

We strongly believe that systematic measurements and data collection are prerequisites for the analysis and management of any service system. Therefore, our students encounter numerous service data sets and examples during lectures and recitations. Most of the course homework is also based on actual service data.

Early generations of the course used one-month tellers’ data from a bank in Israel, in support of recitations and homework. As described in [Mandelbaum and Zeltyn \(1998\)](#), the data were collected through barcodes that were given to customers upon arrival, which were then scanned at milestones of the service experience. Then the focus of the course shifted in the direction of call centers, which seemed to be a natural candidate for becoming our main application area. First, one could not but be impressed by the sheer magnitude of the call center industry, with its explosive growth over the last three decades. Millions of agents are employed in call centers all over the world and, according to some estimates, worldwide expenditure on call centers exceeds \$300 billion per year. Second, the call center environment gives rise to numerous managerial and engineering challenges that vary in their nature and time scale, from real-time skill-based routing calls to long-term design and workforce planning. Thirdly, large call centers generate vast amounts of data that are relatively easily accessible. (For example, hospital data, typically dispersed in a number of different information systems, are more complicated to access.) A detailed history of each call that enters the system can, in theory, be reconstructed via the automatic call distributor (ACD) and interactive voice response units (IVR). However, call centers have not typically stored or analyzed this

data, using instead the ACD reports that summarize performance over certain time intervals (say, 30 minutes). In our research (Brown et al. 2005) and teaching, we advocate the change of this approach and emphasize the practical and research advantages of call-by-call data analysis.

In the ServEng course, our first call-center database covered a one-year operation of a call center at a small Israeli bank (Brown et al. 2005, Mandelbaum et al. 2000): about 350,000 calls, catered by about 15 agents. Once this database was incorporated into the course, the bank tellers' data (face-to-face services) have been used since exclusively in recitations, to teach or demonstrate techniques, while the telephone data have been used in homework to practice these techniques.

Our small Israeli database clearly revealed the potential of incorporating real-world data into ServEng teaching and research. It hence paved the way to access vast amounts of call-by-call data from large call centers that were willing, sometimes even eager, to participate in our data venture. However, then several challenges arose. First, call center data is processed by vendor-specific programs, in formats (data structures) that are typically not amenable to operational analysis. Second, a database of a large call center consists of up to hundreds of millions of records. Hence, brute force statistical processing of the records would take a prohibitively long time, even for the most powerful statistical software. Third, our experience is that a significant part of real-world databases is contaminated or inconsistent with the rest, which calls for an extensive cleaning and data improvement effort prior to using the data.

As one anecdotal example, the first large U.S. bank that provided us with call center data had four interconnected call centers. When switching to daylight savings time, one of these call centers failed to update its computer clock, and thus calls that were dialed in City X arrived at City Y one hour earlier! The database from this U.S. bank, which is described in the US Bank SEELab Report (2006), was in fact cleaned and has recently been incorporated into the ServEng course, as will be described in the next subsection. (Just for comparison sake, the U.S. bank employed about 1,000 call center agents who catered to about 350,000 calls per week—which was the yearly volume of our small Israeli bank.)

### 3.1. DataMOCCA and SEESat—An Environment for Online EDA (Exploratory Data Analysis)

Data-based research and teaching give rise to numerous data-related challenges. These range from partner identification and data acquisition through cleaning to maintenance and analysis. To address the challenges presented by service data, a research laboratory was

established in 2007 at the Technion: SEELab, where SEE stands for “service enterprise engineering” (SEE 2010). The SEELab has been funded by a private donation to the Technion—indeed, research agencies are unlikely to support its activities, as these are not considered basic research (but rather only supporting the latter). Under the dual leadership of Dr. Valery Trofimov (technical) and AM (academic), the SEELab receives, processes, cleans, validates, and maintains a repository of databases from service enterprises. To this end, the SEELab developed DataMOCCA (Data Models for Call Centers Analysis) (Trofimov et al. 2006), which is a software suite that renders the data ready for support of research and teaching.

DataMOCCA offers a *universal* model for operational call center data or, more generally, transaction-based services. (Universality is in the sense of being source-independent, be it a call center or an emergency department or an Internet site.) The main user interface is within a (flexible and friendly) graphical environment, SEESat, which enables real-time exploratory data analysis (EDA), at second-to-month resolutions. As part of the EDA, SEESat includes statistical algorithms such as parametric distribution fitting and selection, fitting of distribution mixtures, survival analysis (mainly for estimating customers' (im)patience), and more—all of these algorithms interact smoothly with all the databases.

The SEELab databases are designed and maintained at two levels: the basic level is close to the raw data, after cleaning, validating, and processing; the upper level consists of precompiled summary tables, which are created only once (and then only occasionally modified), and their design is efficient enough to support real-time processing, with few-seconds response time, covering gigabytes-large databases. This provides the infrastructure for a convenient EDA environment, which accommodates real-time statistical analysis and simulations.

Currently, SEESat interacts with call-by-call data from four large call centers: a U.S. bank, two Israeli banks, and an Israeli cellular-phone company. Three of the four databases cover periods of 2–3 years; the fourth one, which is presently the most active, has about 1.5 years worth of data where, during the last 6 months, call center data has been deposited on a *daily basis* at a data safe within the SEELab.

The U.S. bank database is universally accessible (via the Internet) at the SEELab's server (more on that momentarily). This database has close to 220 million calls, out of which about 40 million were served by agents and the rest by a VRU (voice response unit, or, simply put, an answering machine). DataMOCCA and SEESat have well-served our ServEng course. Indeed, a multitude of examples, using SEESat, are presented in classes and recitations. Moreover, one

of the homework assignments is SEEStat-dedicated: it allows the students hands-on experience with the U.S. bank data mentioned above, having them analyze it in interesting insightful ways. Concrete SEEStat applications from the ServEng course are described in §4 below. Further details, in particular of the SEEStat-based homework, are provided in §§4 and 5 of Mandelbaum and Zeltyn (2009).

While originally designed for operational databases from call centers (hence, its name), DataMOCCA has been expanded to accommodate additional sources and types of data. Specifically, the SEELab now “owns” data also from several hospitals (mostly their emergency departments), from an Internet website (click-stream data), and a couple of samples from face-to-face services. We are further preparing to augment our operational databases with financial/marketing and contents data. (In an emergency department, contents refers to clinical data.) In addition, two simulators have been written, one for a call center, the other for an emergency department. The “vision” is to connect these simulators to the real data of the SEELab and then have these SEELab resources (data and simulators) readily accessible worldwide for use by students and researchers of service engineering and its supporting service science.

**3.1.1. SEEStat Online.** As already mentioned, data from two banks are universally accessible through SEEStat at the SEELab’s server. The connection procedure, for any research or teaching purpose, is simply as follows: go to the SEELab webpage <http://ie.technion.ac.il/labs/serveng>; then, either via the link *SEEStat Online*, or directly through <http://seeserver.iem.technion.ac.il/see-terminal>, complete the registration procedure. Within a day or so, you will receive a confirmation of your registration, plus a password that allows you access to the EDA environment of SEEStat and thus to the abovementioned databases. Note that your confirmation email includes two attachments: a troubleshooting document and a self-taught tutorial. We propose that you print out the tutorial, connect to SEEStat, and then let the tutorial guide you, hands-on, through SEEStat basics—this should take no more than one hour. (For teaching purposes, one of the recitations at the course site contains a SEEStat quick guide, which has been used to introduce it to students.)

## 4. Course Syllabus: Theory, Examples, and Case Studies

The ServEng course consists of roughly four parts:

1. *Prerequisites*: measurements and models;
2. *Building Blocks*: demand, services, customers (im)patience;

3. *Models*: deterministic (Fluid) and stochastic—mainly queueing models, both conventional (Markovian) and approximations; and

4. *Applications*: design, workforce-management (e.g., staffing) and skills-based routing.

We now provide a brief description of these four parts. We continue with a list of the lectures, each accompanied by brief commentary. In our Full Version (Mandelbaum and Zeltyn 2009), readers can find more detailed descriptions of each of the 14 ServEng lectures.

*Measurements*, at the granularity level of individual service transactions, are prerequisites for the design, analysis, and management of service systems. Thus, after opening the course with an introduction to service engineering, we survey transactional measurement systems in face-to-face, telephone, Internet, and transportation systems. We then proceed with an introduction to *Modeling*, using dynamic stochastic PERT/CPM models (also called fork-join or split-match networks) as our modeling framework. These models capture operational congestion that is attributable to resource constraints and synchronization gaps.

Measurements and modeling prerequisites directly give rise to deterministic (fluid/flow) models of a service station, which capture average behavior and enable relatively simple, yet far-reaching analysis—for example, capacity (bottleneck) analysis.

The next course segment is dedicated to the three building blocks of a basic service-model: demand, service, and (im)patience. First we study *service demand*, emphasizing the importance of reliable forecasting techniques. (In particular, our model for exogenous customers’ arrivals is a Poisson process, or a relative). Then we analyze the *service process*, describing its operational characteristics: service-duration, which is a static characteristics, and process structure, capturing dynamic evolution. Service durations in call centers often turn out log-normally distributed (Brown et al. 2005) (at the resolution of seconds, but sometimes exponential at minutes-resolution). The structure of the service process is naturally captured by phase-type distributions. We end with *customers’ patience*, or perhaps *impatience*, and its manifestation—the abandonment phenomena, which is important in call centers and other services (e.g., Internet and even emergency rooms).

The three building blocks, arrivals, services, and (im)patience, are fused into basic queueing models where customers are iid and servers are also iid. A central role is played by Markovian queues, underscoring the applicability of the Erlang-A queue in the call center industry (Mandelbaum and Zeltyn 2007). Then we discuss design principles (pooling to exploit economies of scale) and present operational workforce

management techniques (staffing and scheduling), including staffing in the QED, ED, and QD operational regimes. We conclude the course with models that acknowledge customers differentiation (priorities) and servers heterogeneity/skills (SBR = skills-based routing). An optional last lecture surveys queueing networks, specifically Jackson and generalized Jackson, as models of multistage service systems.

#### 4.1. Course Material and Supporting Texts

ServEng is a one-semester course. Our standard Israeli semester consists of 14 weeks, 3 hours of lectures per week, accompanied by a weekly 1-hour recitation. At every lecture and recitation, students get lecture notes (copies of the slides), which are used in class. Typically, these notes are supplemented with board lectures, providing an introduction to a sub-subject or clarifying subtle issues. (The notes most often constitute a superset of the class material, and the actual material covered varies.) All lecture notes appear in the ServEng website ([Service Engineering 2010](#)), under the lectures and recitations menus.

Israeli students, unlike U.S. students, are not accustomed to using text-books for self-study—they rely on class material, hence, the lecture notes play a central role in the course. Yet, there are three books that are mentioned as relevant and useful, and our library has ample copies of these books for students' use. These books are mentioned at the outset of the course, and students are asked to read few chapters from these books. The books are by [Hall \(1991\)](#), [Fitzsimmons and Fitzsimmons \(2004\)](#), and [Lovelock \(1992\)](#) (in that order of relevance to the course), on which we now elaborate.

Our teaching philosophy was influenced by Hall's book ([Hall 1991](#)) on queueing methods—this book is rather unique among queueing books as it discusses measurements, provides fluid models the respect they deserve, and emphasizes science-based applications. Hall's book serves as an optional textbook, which students are encouraged to consult, especially during the first half of the course. Placing the textbooks of the course in perspective, Hall's book lies at the intersection of operations research and industrial engineering; [Fitzsimmons and Fitzsimmons \(2004\)](#) borders on industrial engineering and operations management; and [Lovelock \(1992\)](#) touches on the four extreme points that span service engineering: operations, marketing, human resource management, and information systems.

Occasionally in our course, we have used a number of business cases, either within course material or as final course projects. An example of course material is the National Cranberry Cooperative case ([Porteus 1989, 1993a, b](#)), used within fluid models. Two final-project examples are Harrison's Manzana Insurance

case ([Harrison et al. 1993](#)), and Federgruen's Vonage case ([Federgruen and Samoednik 2008](#)). (We adapted the latter to our course needs by changing its base model from Erlang-C to Erlang-A.)

We now proceed with listing the topics that we cover, plus some commentary.

**4.1.1. Rules of the Game.** In the first lecture of the course, we introduce students to the service economy and to the discipline of service engineering. We also discuss course logistics and "rules of the game," for example, the relative weights of homework (40%–50%) and final exam (60%–50%) in the final grade.

**4.1.2. Measurements and Models: Little's Law.** The subsequent two lectures are dedicated to *measurements* and *models*—these, in our opinion, are the prerequisites for any advances and practice of engineering and science, in particular to the discipline that we are teaching.

**4.1.3. Data Sources.** As already mentioned, we believe that research and teaching of service engineering must be based on *measurements*: real-world data, collected in various service systems at various levels of granularity. All stages of modeling and design of service systems must be empirically supported, starting with systematic measurements and data collection at the initial stage of modeling and ending with the validation of both models and managerial decisions against empirical data. Call centers provide an excellent illustration for our approach, and their operational data, at the level of individual transaction (*transactional data*), have been our main data source. Recently, we have extended our data sources to cover also hospitals, especially their emergency departments. Additional course data comes from face-to-face, Internet, and transportation services.

**4.1.4. Models.** The second prerequisite for service engineering is *models*. Indeed, a major part of the course is dedicated to different models of service systems, their analysis and areas of application. We distinguish between *empirical* and *analytical* models. (Presently, the course utilizes no *simulation* models, though this is going to change in the near future.) Empirical models, which are directly data-based, are further classified into conceptual, descriptive, and explanatory models. Analytical (or mathematical) models are either deterministic (fluid-based) or stochastic (e.g., Markov chains). All models could be either static (long-run, steady-state) or dynamic (transient). No attempt is made at the course to well-define the abovementioned model classes—these are used merely to structure one's thinking about modeling, and we "explain" them mainly through case studies. For example, some fluid models are in fact

empirical models, as we clarify through Figure 11 in Full Version (Mandelbaum and Zeltyn 2009). Analytical models could overlap with empirical models, as we illustrate via the simple yet fundamental Little's law (Little 1961):  $L = \lambda \times W$ . (The reality of service systems, for example transient behavior, calls for versions and generalizations of Little's law that go beyond steady-state/long-run. We thus pay a significant attention to *time-varying* versions, which includes *finite horizons* as well.)

**4.1.5. Addressing the Skeptic: Is Service Engineering Relevant?** It is not unreasonable to question whether the methods of "exact sciences" (e.g., Mathematics) and their related models (statistical, operations research) are relevant for supporting the operations of the legal system. Indeed, Flanders (1980), who is what we have been calling a *smart influential skeptic*, gives a negative answer to this question. Interestingly, Flanders demonstrates unfamiliarity with Little's law (Flanders 1980, pp. 316–317). We thus believe that the following case study of Little's law, applied to the "production of justice," would have certainly helped a service engineer alter the views of any "Flanders," in favor of the profession.

**4.1.6. Case Study—Little's Law and the "Production" of Justice.** This example, presented in §4.2.4 of Full Version (Mandelbaum and Zeltyn 2009), applies Little's law to data that records the operational performance of five judges in Haifa's labor court. It demonstrates that the *best* operator (judge), having the highest  $\lambda$  and lowest  $W$  (i.e., processing most files with the least delay), has in fact the largest  $L$  (longest queue). Hence, according to the accepted criterion of the justice system, that judge is considered the *worst* operational performer, having the longest file backlog in process.

**4.1.7. Little's Law for Time-Varying Systems.** Toward the end of the course, we teach staffing algorithms for service systems that face time-varying arrival rates (which excludes little else). These algorithms are based on the central notion of *offered load* which, in turn, amounts to a version of Little's law in a *time-varying* environment. We introduce this time-varying version, followed by discussions of the offered load and the relation of the two.

**4.1.8. The Offered-Load.** For a service system in steady state, with constant rate  $\lambda$ , and iid services of (random) durations  $S$ , the *arrival-rate of work* to the system is  $R = \lambda \cdot E[S]$  units of work per unit of time, where work is measured in units of time. We call  $R$  the *offered-load* to the system. What if the arrival rate  $\lambda$  varies in time? The "right" answer conceptualizes the workload offered to a service station in terms of a corresponding infinite-server system, which is natural

and far-reaching. The offered-load is the backbone of time-varying staffing (taught at the end of the course). The concept, which is presently the subject of active research at the Technion, extends far beyond a single basic service station.

**4.1.9. Dynamic Stochastic Service Networks.** In Lecture 4, we formally introduce service networks (or processing networks), which were described previously in §1.3. These provide us with a framework for teaching the main drivers of operational delays: *scarce resources* and *synchronization gaps* (to be distinguished from other causes, beyond the control of the service engineer, for example bad weather in an airport). We actually start with loosely defined conceptual models of service networks, made concrete through call centers and emergency departments. These conceptual models demonstrate (at least) two important points: the power of the "language of modeling" and the role of the service engineer as an *integrator* of several disciplines (recall §1.1).

**4.1.10. Applicable Conceptual Framework.** DS-PERT networks serve as a conceptual framework that teaches students at least the following important points: usefulness of models and the modeling process; clear *process-view* of a service system as a multi-project resource-constraint system (project = customer to be served); and the sources of operational delays, namely, resource-queues and synchronization queues.

**4.1.11. Fluid Models of Service Networks.** Lecture 5 of the ServEng course introduces the fluid approach, which yields the first mathematical models of a service station and a service network. Being intimately related to empirical models, fluid models provide a bridge from the empirical to the theoretical part of our course. Our fluid view is taught through several useful families of models: individual-scenario analysis, queue-buildup diagrams, cumulative flows, and spreadsheet models, the latter based on finite-difference approximations of ordinary-differential-equations (ODEs).

**4.1.12. The Building Blocks of a Basic Service Station.** Starting from Lecture 6, we proceed with a series of lectures on the main building blocks of a stochastic model for a service station. These are: arrival process, representing customer's demand; service process, capturing work per customer that is required from servers; customers (im)patience, which is what distinguishes service to a human customer from that to, say, a part in a production line. Each building block is introduced, both empirically, mainly via SEESat, and theoretically, via a corresponding mathematical model. (After teaching arrivals and services, they are combined to form the workload process.)

**4.1.13. Arrivals of Customers.** In Lecture 6, the students learn about customers' arrivals or, in other words, customers' demand for service. Arrivals to a service system can be studied in several time-scales, each giving rise to a different type of models. In our course, we focus on operational short and middle-term decision making based on arrival-rate patterns, which is well served by Poisson processes—time-homogeneous and nonhomogeneous (the latter modeling daily arrivals in call centers, hospitals, and other service systems).

**4.1.14. The Service Process.** The importance of properly managing the customer-server interaction, or what we call the *service process*, can hardly be underestimated. Lectures 7 and 8 of the ServEng course explore the service process, focusing on its operational attributes: *duration* and *structure* (as opposed to contents or economic worth, for example). Durations are modeled by distributions (empirical, parametric), and structure via phase-type models.

**4.1.15. Customer (Im)Patience.** The subject of customers' (im)patience, e.g., on the phone while waiting to be answered by an agent, in the emergency department while waiting for a nurse or a physician, or on the Internet while waiting for a response, plays a central role in the analysis of service systems. From a practical point of view, (im)patience could lead to customers' *abandonment* which, in our opinion, is perhaps the most important operational performance measure of a call center. (On a personal note, one of the authors (A.M.) is greatly in debt to the abandonment phenomena—it was, in fact, the original trigger of his data-driven approach to service engineering, both research and teaching.) Two important topics that are introduced in the context of (im)patience are censoring and hazard rates, the former because the (im)patience of served customer is left-censored by the actual waiting time. (The actual waiting times of served customers provide lower bounds for their (im)patience.) The latter is introduced because hazard rates constitute natural dynamic models of (im)patience (following Palm 1957, who proposed to use the *hazard rate* of the patience time as a means for *dynamically* describing (im)patience or, rather, customers' irritation as a function of the time of their waiting.)

**4.1.16. Stochastic Models of a Basic Service Station.** Lectures 10–12 of the course are dedicated to models of a basic service station: with a homogeneous customers' population (as opposed to multitype) and iid servers (as opposed to varying skills). Specifically, Lecture 10 provides a survey of the classical Erlang-C ( $M/M/n$  model), and it introduces the 4CallCenters software (4CallCenters Software 2002), or 4CC for short—a user friendly, effective practical tool for the analysis of basic Markovian queues (Erlang A/B/C

and relatives). Lecture 11 reiterates the importance of the abandonment phenomena and presents the basics of Erlang-A ( $M/M/n+M$ )—the model which, in practice, is gradually replacing Erlang-C as the queueing engine of workforce management software. Finally, Lecture 12 covers non-Markovian queues, with general service and interarrival distributions ( $G/G/n$ ).

**4.1.17. Operational Regimes and Staffing—QED Queues.** A central challenge in the design and management of a service operation in general and of a call center in particular is to achieve a balance between *operational efficiency* and *service quality*. This is addressed in Lecture 13, where we show that, depending on the desired balance, several appropriate operational regimes could arise through corresponding staffing rules. Among these, we introduce the ED, QD, and QED regimes.

**4.1.18. The QED Regime.** The QED regime goes hand in hand with square-root staffing:  $n = R + \beta\sqrt{R} + o(\sqrt{R})$ , where  $R$  is the offered-load, and the service grade  $\beta$  is determined by the (limiting) delay probability, say,  $\alpha$  (and vice versa). The function  $\beta(\alpha)$ , for Erlang-C, is called the Halfin-Whitt function. For Erlang-A, it is called the Garnett function, and it is computable via a spreadsheet. QED approximations are typically extremely accurate over a wide range of parameter values from the very large call centers (1,000s of agents) to the moderate-size (few 10s of agents) (Janssen et al. 2010); this renders QED approximations applicable to small systems, such as those found in healthcare.

**4.1.19. Time-Stable Performance of Time-Varying Queues.** The staffing rules based on Erlang-B/C/A assume that all system parameters, in particular the offered load, do not vary with time. This is definitely not the case in call centers and most other service systems, which are clearly time inhomogeneous. Performance analysis of time inhomogeneous systems can be addressed within the framework developed in Feldman et al. (2008), where traditional staffing approaches are rephrased as follows: Given a time-varying arrival process, what is the staffing level (necessarily, time-varying) under which operational performance is time-stable? In other words, Feldman et al. (2008) answers the question of how to time-stabilize the performance of a time-varying system. In the answer, the time-varying Little's law/offered-load play a central role.

**4.1.20. Workforce Management: Hierarchical Operational View.** While we teach mainly staffing methods within operational WorkForce Management (WFM), we do emphasize that the spectrum of WFM is far broader. The hierarchy of decisions along the WFM chain goes through the following steps: forecasting, staffing, shifts, rostering, and skills-based

routing. This hierarchy is not unique to call centers (e.g., nurse staffing in hospitals). Except for rostering, all steps of the hierarchy are covered in either lectures or recitations, with students learning and practicing forecasting, queueing models, shift scheduling (solving an IP), and some SBR principles.

**4.1.21. Heterogeneous Customers and Servers: Skills-Based Routing (SBR).** Our models have assumed, so far, iid customers and iid servers. A next step toward the complex reality of service systems is to allow heterogeneity of either customers or servers, or both—which amounts to studying *queueing networks* (recall §1.3.) To maintain some level of analytical tractability, one can proceed in one of two ways: Either render complex the topology of the network while maintaining its protocols simple, for example, as in Jackson networks; or, conversely, assume a simple network topology while allowing complex protocols, as is the case in SBR or Fork-Join networks. In view of time and scope constraints, only one of these models can be covered in ServEng—and we chose SBR: it seemed to be the more natural path when one is teaching services, and it was consistent with the course’s emphasis on call centers. Thus, the last lecture covers SBR. We have been taking a practically oriented approach, covering parts of the teaching note of Garnett and Mandelbaum (2000), and expanding with some industry papers and recent research results when feasible, and as time permits.

REMARKS.

- Given the wide spectrum of models, covered in our course, a legitimate question can be asked: How can one choose an appropriate model for a specific practical problem? Of course, no universal recommendation can be provided, because model selection can be an art, as well as a science. As a general guideline, we suggest to start from the simpler models (fluid approximations, for example) and proceed to more complex models if the simpler ones are inapplicable. Recall also the four questions considered in §1.3.2. Very often, the complexity of the relevant model increases as one progresses from the first question (“Can we do it?”) to the last one (“What is optimal to do?”).

- It is clearly impossible to cover all important service engineering methods and applications within a single-semester course. Moreover, our topic selection was much a consequence of personal taste, data availability and course-time constraints. Thus, a service engineering course that is taught elsewhere, by another teacher, could legitimately have another application focus and include significantly alternative material.

- Service sectors that could be alternative to call centers are healthcare (which, in fact, we have been

gradually incorporating into our course), government services, financial services, Internet services, and more. Additional or alternative topics, relative to what we have been covering, could be topics on the interface of operations with other fields (marketing, human resource management, information systems), yield management, richer queueing models (e.g., systems with appointments, various networks), extended theoretical focus (statistics, operations research), service quality, service location, service automation, etc.

- Until recently, an evaluation form was distributed to students during one of the last classes. Students then provided their course assessment, in terms of many parameters. In recent years, the paper form has been replaced by an Internet questionnaire. As a rule, the “course grades” have been very good to excellent.

## 5. ServEng Homework and Exams: A Data-Based Approach

In concert with the way the course is taught, both in lectures or recitations, the course assignments and exams are also data based, when feasible. ServEng students are thus exposed to various real-world data sets, with which they apply theoretical and practical tools that they were taught. The data-based exercises still go hand in hand with ample interesting theoretical homework, so that understanding the material goes beyond the “recipe level.” Homework weighs about 50% towards the final grade, while the rest is determined by a final exam (3–3.5 hours). In this section we describe both homework and exams.

Data-based homeworks and exams are difficult to create, hence the support of the SEELab is essential here, and, even so, the homeworks must serve the course over a relatively long time, and writing an exam takes many hours and iterations (a joint Instructor + TA effort).

Data analysis tends to require relatively many hours of work from students, some of it not directly rewarding, yet it is very important for the learning experience of ServEng. Hence homeworks are carried out in groups (up to four students in earlier generations of the course and student pairs in recent generations), and their weight in the final grade is substantial—around 50%, as already mentioned. Group assignments also save on feedback and grading resources, which are not in abundance. (High-quality feedback is a must for the hard-working students, to help them both learn from their mistakes as well as sustain their motivation.) To achieve all of the above, the support staff of the ServEng course (TAs, homework graders) must be dedicated and of the highest quality, which has indeed been the case—we return to describing their essential contributions in §6.

## 5.1. Homework Assignments

We now present a complete list of the homework assignments, with accompanying brief descriptions.

**5.1.1. List of Assignments.** Recitations and assignments, as well as additional related materials, can be downloaded from the ServEng website. Note that one can also download partial solutions for most of the assignments—these were installed when we moved from groups of four to pairs, to reduce students' workload. (Professors who would like to experiment with teaching ServEng could write to the first author (A.M.), regarding further information about the homework.)

1. **On Queues.** This introductory assignment requires from students to read some textbook/background materials and present a brief executive summary on "The Future of Queues in Service Systems."

2. **Capacity Analysis. Little's Law.** The assignment contains mostly theoretical questions on Little's law; some were in fact motivated by empirical scenarios.

3. **Empirical Models.** This is the first data-based homework assignment. Data of a U.S. banking call center, exported via DataMOCCA (Trofimov et al. 2006), is prepared as convenient Excel files that students then use. Answering the questions requires the application of lecture material on Little's law, capacity analysis, and fluid models.

4. **Processing Networks.** We ask students to create a model of a process network (see §4.3 of Full Version (Mandelbaum and Zeltyn 2009)) with which they are familiar. First, an activity precedence diagram should be described, then a resource-based representation should be given and, last, the two representations must be integrated into the combined diagram of the activities and resources. An information flow diagram has been optional in recent years. In addition, students are required to speculate on the service engineering of their model, based on their diagrams.

5. **SEESat Analysis.** Students download an educational version of SEESat (the GUI of DataMOCCA), with which they explore part of the U.S. Bank database. They search for answers to several practically oriented questions, redoing (at the push of a button) some of the exercises that they worked hard over, in the previous empirical models homework. We elaborate on this homework in §5.1.2 of Full Version (Mandelbaum and Zeltyn 2009).

6. **Fluid Models.** The students address several questions, related to call center applications, via the fluid approach: they write the proper differential equations, solve them in Excel using the finite difference method, and finally apply the Excel Solver to optimize some aspect of the operation (e.g., SBR, overtime management). See §5.1.2 of Full Version (Mandelbaum and Zeltyn 2009) for the detailed coverage of this homework.

7. **Statistical Analysis of Arrivals.** This homework includes some statistical question (e.g., practical testing of the Poisson hypothesis) and introductory-level forecasting problems.

8. **Service Processes and Empirical Analysis of Customers' (Im)Patience.** Some questions here are purely theoretical and others, on abandonment behavior, require the interpretation of real-data from call center.

9. **Gazolco's Call Center.** As mentioned, 4CallCenters (4CC) is software (4CallCenters Software 2002) that supports the application of various Markovian queueing models (Erlang-C, Erlang-B, Erlang-A, and others). 4CC is a valuable tool, serving as a personal WFM tool that students actually use also after graduation. It is practiced via operational problems of the fictitious Gazolco call center, which is an assignment that helps students develop their intuition on practical call center challenges.

10. **Theoretical Analysis of Service Stations in Steady State: Priority Queues.** This homework includes a small case study on pizza delivery, adapted from Fitzsimmons and Fitzsimmons (2004). It is followed by some queueing-theoretical questions.

11. **Staffing of Call Centers.** This last homework (in essence, the final project) is based on actual call center data (we have used various call centers across different semesters). It asks the students to apply a variety of tools they acquired during ServEng lectures. In one version of the homework, data from three call centers was used: a small Israeli call center, for applying standard queueing models; a large U.S. call center that calls for QED analysis; and a medium Italian call center, which went through an interesting process of "hiring and firing."

## 5.2. The Final Exam

Until several years ago, exams were very comprehensive, requiring a labor-intensive preparatory process. All exams + solutions appear on our website, but they are written in Hebrew. An English-version example of such an exam, with its solution, appears in Mandelbaum (2004).

The second generation of our exams is of a different form, which is more structured and hence easier to prepare. The exam consists of four questions (for a total of around 50 points = the weight of the exam in the final grade):

1. **Homework Question:** Taken from the homework and worth about 10–15 points. This question verifies that students were not free-riders when preparing their group homework.

2. **Recitations + Lectures Question:** Taken from lecture-notes provided in recitations and lectures and worth about 10–15 points. This question encourages class attendance, which enhances the learning experience (and which students too often tend to neglect).

3. **Practical Question:** Asking the students to actually analyze data, e.g., identifying operational regimes and some of their properties, which is worth about 15 points. (The students do not need a computer to work on this question: it is composed to be solved with paper and pencil.) Because in their future work most students will encounter real data, this question in fact tests what they are actually carrying with them after graduation.

4. **Theoretical Question:** Asks students to prove or interpret part of a theoretical fact (e.g., Khintchine-Pollaczek, or biased sampling from PASTA, or something similar to a proof from class), worth about 10 points. This question is a prerequisite for getting one's final grade up to the level of 90–100. Students who perform well in this question often continue to graduate school (or are already there).

## 6. Acknowledgments and a Little More History

The present service engineering course at the Technion is the culmination of a long continuous-improvement process. One can view this process as progressing along several fronts:

- *Teaching material:* Most of the material is original, based on research papers of the authors and coauthors, customized teaching notes, graduate theses, and students' projects.

- *Recitations and Homework:* These have been developed originally by the authors and then enhanced, expanded, and innovated by the many excellent teaching assistants (TAs) that the course and its students have enjoyed. As already mentioned, homeworks are central to the learning experience, which is manifested by a weight of about 40%–50% of the final grade.

- *Students' Projects:* Undergraduate IE&M student at the Technion must take part in a yearly project, typically carried out during one's last (4th) year of school, with the goal of applying theory to a real-world setting. Similarly, graduate students in the IE&M Operations Research program take a project-course, with the same goal. A full list of projects that the first author has advised appears in [http://ie.technion.ac.il/serveng/references/avim\\_cv\\_08.pdf](http://ie.technion.ac.il/serveng/references/avim_cv_08.pdf). Many of these projects apply tools acquired in the service engineering course (mostly to call centers and hospitals)—then excerpts from such projects are often used to demonstrate service engineering applications. Other projects constitute pilots that develop into research projects or graduate theses.

For example, Figures 7–10 in Full Version describe the routing process of hospital patients from an emergency department to internal wards. These were first generated within an undergraduate project at the Rambam Hospital in Haifa, then continued as the

M.Sc. thesis of Tseytlin (2009) and used as an example for the homework where students are asked to create their own processing network. As another example, Figure 20 in Full Version was first created within a project at a call center in a large Israeli Bank. It has become part of a research paper that is now in the writing.

- *Website:* Maintenance of the course's website has been the responsibility of the course's TA. While there is a high toll in terms of time and effort, the website has been a central enabler for development and maintenance. Indeed, the website serves as an organized repository of course material, active and archival, for the benefits of students, instructor, and TAs. Also, through the website, course material is accessible to teachers and research colleagues worldwide (which is the reason that the material has been developed in English, while the Technion course has been taught in Hebrew).

*Course TAs:* The service engineering course owes much of its success to its excellent TAs. They have been carefully chosen and then well-groomed for the job (for example, by reattending lectures as a future TA and learning "trade secrets" from the present TA). Our TAs have all been top graduate students, doing research directly related to some course subjects and willing to undertake the challenging job of a ServEng TA. This entails first the routine yet time-consuming chores of conducting recitations, maintaining the course website, having office hours (which are usually attended because of the challenging homework) and, all in all, help the course maintain a service level that a service engineering course is worthy of. In addition, there is also the creative part of a TA work (which distinguishes the excellent from the merely very good)—that of generating new material for recitations and homework and writing questions for exams, with an emphasis given to incorporating the TA's own research into course material. One of the authors (S.Z.) was the first TA for the course in its present form. Then, in chronological order, we acknowledge the significant contributions of TAs Itay Gurvich, Gennady Shaikhet, Shimrit Maman, and Galit Yom-Tov.

As for *student contribution*, there is a long list of (mostly graduate) students who have contributed to the course material development (sometimes while serving as homework graders). Among these, we would especially acknowledge Sivan Aldor-Noiman, Yonit Baron, Izik Cohen, Ofer Garnett, Zohar Feldman, Eva Ishay, Polina Khudyakov, Pablo Liberman, Yariv Marmor, Michael Reich, Luba Rosensmidt, Arik Senderovich, Yulia Tzeytlin, and Asaf Zviran. As for *The SEELab*, permanent and temporary members of the SEELab (SEE 2010) have supported the course in various ways, mostly related to data and

SEESat. In this regard we acknowledge Valery Trofimov, Ella Nadjarov, Igor Gavako, Katya Kutsy, and Arik Senderovich.

As for *research partners*, when teaching, especially a young subject, it is useful for teachers to reflect and students to study the subject's evolution. This is also an opportunity to acknowledge those who contributed to that evolution. An example is the overhead in Figure 2, taken from a lecture on queues with impatient customers: it describes the "modeling history" of such queues, as inspired by call centers; it also lists the models and their developers, who have been our coauthors and partners and to whom we are highly grateful. Names of graduate students are highlighted in the slide to stimulate the curiosity and appetite of course students toward projects and graduate studies in the field of service engineering.

In terms of *administration support*: Special thanks are due to the current dean of the IE&M Faculty at the Technion, Professor Boaz Golany, who followed closely the development of the course and spared no resources to nurture it. In terms of *minicourses*: Miniversions of the course have been taught at several universities: Wharton (hosted by Larry Brown, Noah Gans, and Morris Cohen), Columbia GSB (hosted by Awi Federgruen and Ward Whitt), Stanford GSB (hosted by Mike Harrison and Sunil Kumar), and INSEAD (hosted by Zeynep Aksin). These minicourses gave us the opportunity to share and learn

from our hosts and their students—their generosity, along all dimensions, is greatly appreciated.

It is important to note that *subsequent courses to service engineering*, advanced service engineering courses are occasionally taught at the Technion (though not at the frequency we desire, because of scarce resources). They either focus on a specific application (a recent seminar was dedicated to healthcare) or on advanced theoretical techniques (e.g., fluid and diffusion approximations of queueing systems). In addition, the first author (A.M.) is presently leading a multidisciplinary seminar in services, attended by both faculty and (mostly honor) students. (IBM provided summer scholarships for some of the students, who were engaged in preparing the seminar.) In this seminar, faculty from various disciplines, e.g., marketing and psychology, will lecture on "services" from the viewpoint of their discipline. The seminar is then planned on giving rise to multidisciplinary undergraduate projects: for example, jointly Operations and Information System (on RFID-based measurements in hospitals), Operations and Psychology (on actual and perceived workload in services), and more. These projects will be jointly supervised by faculty from the corresponding disciplines.

Finally, regarding "*customers*" or "*products*": Is the process of teaching service engineering a production process (contributing to the production of service engineers) or a service process (serving students)? This is a question worth discussing in a first lecture of a service engineering course, but we shall not pursue it here. We shall only say that, as either "service-providers" or "producers," in short teachers, we are grateful to the many students of service engineering, who have learned and labored through the course material and homework: undergraduates, graduates, postdocs, researchers, practitioners, and colleagues. The Technion, and the abovementioned hosting institutions for the minicourses, are blessed with students of high quality and drive, which makes the teaching experience (more often than not) and, hopefully, the learning experience as well, challenging, enjoyable and rewarding. Indeed, learning that a course graduate keeps the lecture notes handy on an office shelf, or having the course be the trigger for a student's successful academic career, is all that a teacher can hope for.

**Figure 2 Acknowledgment to Partners, Coauthors, and Students**

**Call Centers = Q'sw/Impatient Customers  
14 Years History, or "A Modelling Gallery"**

1. Kella, Meilijson: Practice  $\Rightarrow$  Abandonment important
2. Shimkin, Zohar: No data  $\Rightarrow$  Rational patience in equilibrium
3. Carmon, Zakay: Cost of waiting  $\Rightarrow$  Psychological models
4. Garnett, Reiman, Zeltyn: Palm/Erlang-A to replace Erlang-C/B as the standard steady-state model
5. Massey, Reiman, Rider, Stolyar: Predictable variability  $\Rightarrow$  Fluid models, Diffusion refinements
6. Ritov, Sakov, Zeltyn: Finally data  $\Rightarrow$  Empirical models
7. Brown, Gans, Haipeng, Zhao: Statistics  $\Rightarrow$  Queueing science
8. Atar, Reiman, Shaikhet: Skills-based routing  $\Rightarrow$  Control models
9. Nakibly, Meilijson, Pollatchek: Prediction of waiting  $\Rightarrow$  Online models and real-time simulation
10. Garnett: Practice  $\Rightarrow$  4CallCenters.com
11. Zeltyn: Queueing science  $\Rightarrow$  Empirically based theory
12. Borst, Reiman; Zeltyn: Dimensioning  $M/M/N+G$
13. Kaspi, Ramanan: Measure-valued models and approximations
14. Jennings; Feldman, Massey, Whitt: Time-stable performance (ISA)

## Supplementary Material

Files that accompany this paper can be found and downloaded from <http://ited.pubs.informs.org>.

## References

- Borst, S., A. Mandelbaum, M. Reiman. 2004. Dimensioning large call centers. *Oper. Res.* 52(1) 17–34.

- Brown, L. D., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: A queueing science perspective. *J. Amer. Statist. Assoc.* **100**(469) 36–50.
- 4CallCenters Software. 2002. 4Call Centers Software. <http://ie.technion.ac.il/serveng/4callcenters/downloads.htm>.
- Corbett, C. J., L. N. Van Wassenhove. 1993. The natural drift: What happened to operations research? *Oper. Res.* **41** 625–640.
- Erlang, A. K. 1948. On the rational determination of the number of circuits. E. Brockmeyer, H. L. Halstrom, A. Jensen, eds. *The Life and Works of A. K. Erlang*. The Copenhagen Telephone Company, Copenhagen.
- Federgruen, A., A. Samojednik. 2008. VONAGE: Designing a service strategy. Columbia University Case, New York. <http://www.docstoc.com/docs/14384169/vonage>.
- Feldman, Z. 2008. Optimal staffing of systems with skills-based-routing. Technion M.Sc. thesis. Haifa, Israel. [http://ie.technion.ac.il/serveng/references/zohar\\_thesis.pdf](http://ie.technion.ac.il/serveng/references/zohar_thesis.pdf).
- Feldman, Z., A. Mandelbaum, W. Massey, W. Whitt. 2008. Staffing of time-varying queues to achieve time-stable performance. *Management Sci.* **54**(2) 324–338.
- Fitzsimmons, J., M. Fitzsimmons. 2004. *Service Management: Operations, Strategy, Information Technology*, 4th ed. McGraw-Hill, New York.
- Flanders, S. 1980. Modeling court delay. *Law Policy Quart.* **2** 305–320.
- Frei, F. X., P. T. Harker, L. W. Hunter. 1998. Innovation in retail banking. Report, Wharton's Financial Institutions Center. <http://ie.technion.ac.il/serveng/lectures/retail.pdf> or <http://fic.wharton.upenn.edu/fic/papers/97/9748.pdf>.
- Garnett, O., A. Mandelbaum. 2000. An introduction to skills-based routing and its operational complexities. Teaching note, Technion, Haifa, Israel. <http://ie.technion.ac.il/serveng/lectures/sbr.pdf>.
- Garnett, O., A. Mandelbaum, M. Reiman. 2002. Designing a telephone call-center with impatient customers. *Manufacturing Service Oper. Management* **4** 208–227.
- Halfin, S., W. Whitt. 1981. Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* **29** 567–588.
- Hall, R. W. 1991. *Queueing Methods for Services and Manufacturing*. Prentice-Hall, Englewood Cliffs, NJ.
- Hall, R. W., ed. 2006. *Patient Flow: Reducing Delay in Healthcare Delivery*, Vol. 91. International Series in Operation Research and Management Science, Springer, New York.
- Harrison, M. J., C. Loch, D. P. Grant. 1993. Manzana insurance: Fruitvale branch. Stanford University Graduate School of Business School Case S-DS-87, Palo Alto, CA.
- Herman, R. 1992. Technology, human interaction, and complexity: Reflection on vehicular traffic science. *Oper. Res.* **40**(2) 199–212.
- Janssen, A. J. E. M., J. S. H. van Leeuwen, B. Zwart. 2010. Refining square root safety staffing by expanding Erlang C. *Oper. Res.* Forthcoming.
- Jennings, O., F. de Vericourt. 2007. Nurse-to-patient ratios in hospital staffing: A queueing perspective. Working paper, Duke University, Durham, NC.
- Khudyakov, P. 2006. Designing a call center with an IVR (interactive voice response). M.Sc. thesis, Technion, Haifa, Israel. [http://ie.technion.ac.il/serveng/references/thesis\\_polyna.pdf](http://ie.technion.ac.il/serveng/references/thesis_polyna.pdf).
- Little, J. D. C. 1961. A proof of the queueing formula  $L = \lambda W$ . *Oper. Res.* **9** 383–387.
- Lovelock, C. G. 1992. *Managing Services: Marketing, Operations and Human Resources*. Prentice-Hall, Englewood Cliffs, NJ.
- Mandelbaum, A. 2004. Example of a ServEng exam. [http://ie.technion.ac.il/serveng/lectures/exams/moed\\_a\\_2004\\_eng.pdf](http://ie.technion.ac.il/serveng/lectures/exams/moed_a_2004_eng.pdf).
- Mandelbaum, A. 2010. Curriculum vitae. Accessed September 5, 2010, [http://ie.technion.ac.il/serveng/references/avim\\_cv\\_2010.pdf](http://ie.technion.ac.il/serveng/references/avim_cv_2010.pdf).
- Mandelbaum, A., S. Zeltyn. 1998. Estimating characteristics of queueing networks using transactional data. *Queueing Systems: Theory Appl. (QUESTA)* **29** 75–127.
- Mandelbaum, A., S. Zeltyn. 2007. Service engineering in action: The Palm/Erlang-A queue, with applications to call centers. D. Spath, K.-P. Fähnrich, eds. *Advances in Services Innovations*. Springer-Verlag, Berlin, 17–48.
- Mandelbaum, A., S. Zeltyn. 2009. Service engineering: Data-based course development and teaching. Full Version. (A full version of the present document will appear on the *INFORMS Transactions on Education* website.)
- Mandelbaum, A., W. A. Massey, M. Reiman. 1998. Strong approximations for Markovian service networks. *Queueing Systems: Theory Appl. (QUESTA)* **30** 149–201.
- Mandelbaum, A., A. Sakov, S. Zeltyn. 2000. Empirical analysis of a call center. Technical report, Technion, Haifa, Israel. <http://ie.technion.ac.il/serveng/references/ccdata.pdf>.
- Newell, G. F. 1982. *Application of Queueing Theory*. Chapman and Hall, London.
- Palm, C. 1957. Research on telephone traffic carried by full availability groups? *Tele.* (English translation of results first published in 1946 in Swedish in the same journal, which was then entitled *Tekniska Meddelanden Fran Kungl. Telegrafstyrelsen* **1** 107.)
- Porteus, E. L. 1989. The case analysis section: National Cranberry Cooperative. *Interfaces* **19**(6) 29–39.
- Porteus, E. L. 1993a. Case analysis: Analyses of the National Cranberry Cooperative: 1. Tactical options. *Interfaces* **23** 21–39.
- Porteus, E. L. 1993b. Case analysis: Analyses of the National Cranberry Cooperative: 2. Environmental changes and implementation. *Interfaces* **23** 81–92.
- SEE: Technion Service Enterprise Engineering Research Center, Haifa, Israel. Accessed September 5, 2010. <http://ie.technion.ac.il/labs/serveng/>.
- Service Engineering. 2010. Technion, Haifa, Israel. Accessed September 5, 2010. <http://ie.technion.ac.il/serveng>.
- Simon, H. A. 1996. *The Sciences of the Artificial*, 3rd ed. The MIT Press, Cambridge, MA.
- Trofimov, V., P. D. Feigin, A. Mandelbaum, E. Ishay, E. Nadjarov. 2006. DATA model for call center analysis: Model description and introduction to user interface. Technical Report, Technion–Israel Institute of Technology, Haifa, Israel <http://ie.technion.ac.il/labs/serveng>.
- Tseytlin, Y. 2009. Queueing systems with heterogeneous servers: On fair routing of patients in emergency departments. M.Sc. thesis, Technion, Haifa, Israel. <http://ie.technion.ac.il/serveng/references/thesis-yulia.pdf>.
- US Bank SEELab Report. 2006. DataMOCCA: DATA model for call center analysis—The call center of US Bank. [http://ie.technion.ac.il/labs/serveng/files/the\\_call\\_center\\_of\\_us\\_bank.pdf](http://ie.technion.ac.il/labs/serveng/files/the_call_center_of_us_bank.pdf). (Prepared by Olga Donin, with the support of SEELab researchers.)
- Whitt, W. 2010. Research site. Accessed September 5, 2010. <http://www.columbia.edu/~ww2040/allpapers.html>.
- Yom-Tov, G. 2007. Queues in hospitals: Semi-open queueing networks in the QED regime. Ph.D. Research Proposal, Technion, Haifa, Israel. [http://ie.technion.ac.il/serveng/references/proposal\\_galit.pdf](http://ie.technion.ac.il/serveng/references/proposal_galit.pdf).
- Zeltyn, S., A. Mandelbaum. 2005. Call centers with impatient customers: Many-server asymptotics of the  $M/M/n + G$  queue. *Queueing Systems: Theory Appl. (QUESTA)* **51** 361–402.
- Zviran, A. 2008. Fork-join networks in heavy traffic: Diffusion approximations and control. M.Sc. Research Proposal, Technion, Haifa, Israel. [http://ie.technion.ac.il/serveng/references/asaf\\_research\\_proposal.pdf](http://ie.technion.ac.il/serveng/references/asaf_research_proposal.pdf).