Service Comprehension

Richard Taylor, Chris Tofts

HP Laboratories
HPL-2008-140

Keyword(s):
services, Systems modelling, concurrency, comprehension

Abstract:

Confronted with the complexity of services systems, providers capitulate to the *build and hope* approach, so dominant in software engineering - rather than the design, analyze, build and improve methods of mature, science backed, engineering disciplines. Providers simply assume that such systems cannot be predicted, understood or forecasted! Given the increasing dependence on advanced economies on such complex services we need methods of analysis and comprehension so that all of the stakeholders can understand them. Accepting that this cannot take place in a context of a 'professional modeler' approach, which leads to only the modeler owning the understanding. Equally, customers as stakeholders of complex services must move from a position of demand without comprehension of consequences if they are to receive genuinely beneficial services. We must be clear how we exploit the location of boundaries in order that we can comprehend the services we desire whilst understanding the impact that in many service encounters the customer, often the citizen, does not provide direct economic feedback. Finally perhaps, we should acknowledge that many of the tools required could already exist - it is simply that the various stakeholders in a service provision encounter are unaware or intimidated by them.

External Posting Date: October 21, 2008 [Fulltext] Approved for External Publication
Internal Posting Date: October 21, 2008 [Fulltext]

Service Comprehension

Chris Tofts & Richard Taylor,
HP Laboratories Bristol
{chris.tofts, richard.taylor}@hp.com

If we designed cars as well as we design services, then they would have one axle and five wheels. Behara and Chase (1993)

Abstract

Confronted with the complexity of services systems, providers capitulate to the 'build and hope' approach, so dominant in software engineering - rather than the design, analyze, build and improve methods of mature, science backed, engineering disciplines. Providers simply assume that such systems cannot be predicted, understood or forecasted! Given the increasing dependence on advanced economies on such complex services we need methods of analysis and comprehension so that all of the stakeholders can **understand** them. Accepting that this cannot take place in a context of a `professional modeler' approach, which leads to only the modeler owning the understanding. Equally, customers as stakeholders of complex services must move from a position of demand without comprehension of consequences if they are to receive genuinely beneficial services. We must be clear how we exploit the location of boundaries in order that we can comprehend the services we desire whilst understanding the impact that in many service encounters the customer, often the citizen, does not provide direct economic feedback. Finally perhaps, we should acknowledge that many of the tools required could already exist - it is simply that the various stakeholders in a service provision encounter are unaware or intimidated by them.

Why do we need to understand services?

Within the world's advanced economies services are now the dominant element, accounting for approximately 80% of economic activity. However, much as Adam Smith identified services as non-productive activities there is a continuing problem with improving (or indeed measuring) the performance of service delivery. One of the main reasons why the service element of the economy has become so dominant could well be that almost everything can be considered as a service, from the perspective of comprehension the inability to distinguish the object understudy from what should not be understudy is a major challenge. Studies, (Samson, S., 2007) show that there is almost no viable definition of service which would achieve the distinction between activities that one would seek. Unlike the manufactured object there is almost no point at which the boundary of the service can be identified. Indeed there is the temptation when dealing with understanding services to permit the boundary of the study to drift wider and wider. Inevitably when the object understudy is addressable from many points of view, the natural temptation of any expert is to apply their own point of view to greatest effect, tending to avoid the value obtained from multiple points of view.

In the business to business context some of the service offerings are phenomenally complicated, they can involve the use of very advanced technologies, often spread over multiple and disparate technological domains, and with a part of the service offering usually being directly provided by highly skilled human beings. This combination of large amounts of technical systems both hardware and software, and highly trained humans, means that understanding what and how the beneficial service from the system will be delivered is itself an immensely challenging task. It is not uncommon in the information technology services space for both the number of computers

and the number of staff to be in the thousands. Equally, the number of potentially distinguishable services which are supported by these human-technical systems can often be numbered in the hundreds. These systems encompass almost all of the elements which individually place major demands on our ability to comprehend, in particular they exhibit: complicated interactions (many different forms of communications, machine to machine, machine to human, human to human); widely heterogeneous technologies; large scale software systems and the unavoidable presence of humans. Faced with these challenges it is hardly surprising that the primary response is simply to deploy the system, observe its behaviour and hope for inspiration to resolve any particularly poor behaviour.

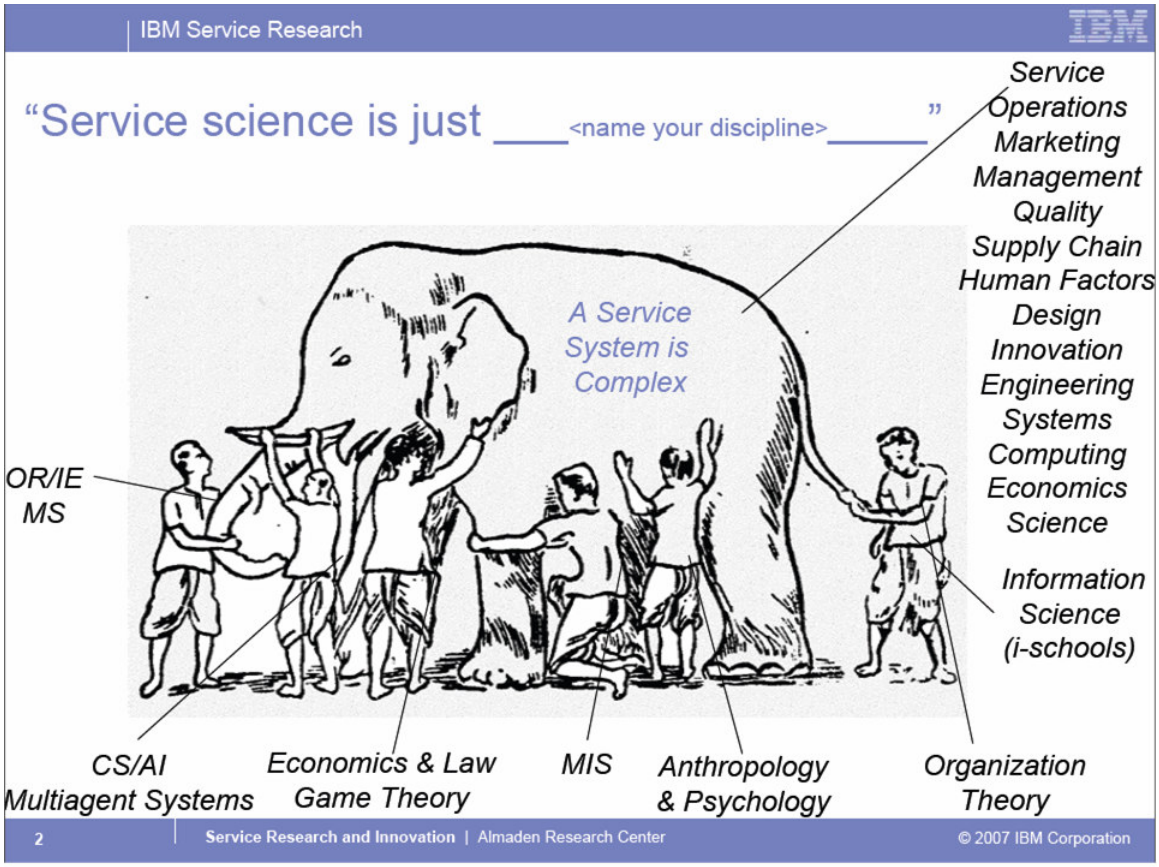

Figure 1 Jim Spohrer's elephant from Spohrer (2006).

Historically, faced with challenges of the scale above (such as behaviour of natural populations, properties of gasses, multiple bodies in orbits, behaviour of fluids,…), the natural approach has been to abstract. However, given the interdisciplinary nature (See Figure 1) of the primary knowledge which drives the design of service systems, there is a competition between the various disciplines as to what constitutes a reasonable abstraction. Different activities have widely different views on the consequential risks of the omissions of description forced by particular choices of abstraction. Indeed given that humans are correctly identified as primary actors within the systems there are many who question whether abstraction can be undertaken at all. One of the primary reasons to avoid abstracting humans in the industrial context is that it is widely associated with Ford-ist views of management and the belief that this can not lead to an appropriate treatment of highly trained and contingent working staff. Unfortunately, if we approach the human aspects of service delivery with a view that they cannot and must not be measured, this leaves us with little ability to improve the system that they sit within. Whilst inevitably there is a view that service system improvement is largely to remove the costly human elements, it is actually possible with correct measurement that the improvements could deliver

both an improved working environment for the human and the more economically viable service for both the provider and the consumer. But, in the absence of an understanding of the interaction of the human behaviour and the service delivery, neither of these benefits is likely to accrue and the inevitable focus of service change will be to reduce cost.

A major challenge facing any organisation attempting to understand a service is the bewildering range of potential approaches to the comprehension task. This form of "target overload" often leads to no successful activity being undertaken at all - similar to how herd animals avoid predators. The decision to commit to, and difficulty of the comprehension task, is often made more difficult by the attempt to perform it all in one place or methodology. A particularly common form of this is the "master simulation", that is the entirety of the problem captured by an executable form of software representation. The usual consequence of this approach is that the only person who understands the system representation is the modeller hired to generate the model. Whilst it appears that we can emulate any system, no matter how complex, within silicon. It is entirely unclear what the value of such emulation is. Are we simply replacing experimentation on the final delivered system with experimentation on the in silico system. Indeed paraphrasing Alan Bundy (1988) have we simply taken *"one system that we do not understand and replace it with another which we do not understand"*. Finally, such approaches are very difficult to maintain, since their complexity matches that of the service system the difficulty of maintaining them is similar. We need to be clear what the intent of the scope (Jain, 1991, pg 22) of a service comprehension exercise are, in particular who owns the comprehension project and who is intended to use and benefit from its output. If neither of these parties can understand either the comprehension approach or the comprehension outputs then it is very hard to see that it has any value.

Even within a single organisation there is a major requirement for service comprehension. It is unusual for any service to be delivered by a single component of an organisation, even in the simplest example there is usually a separation between the manufacturing and the servicing of a classic product. Different parts of an organisation will often have entirely different approaches to how they understand the problem of delivering their element of a service. Furthermore, they will often have different measures of their performance; in particular approaches to revenue recognition can vary dramatically within an organisation. In this setting each part of the organisation is likely to endeavour to optimise their part of the delivery, however unless the service (for example delivery of computation by PC) is designed to exploit the local optimisation efficiently this is unlikely to lead to observable improvement in the overall service delivery. In many cases the competition between elements of the organisation, to be permitted to optimise their parts of the delivery, may well have an extremely negative impact on the other parts of the organisation. For example highly customisable Web services may be particularly attractive to a marketing function, but can place an unacceptable load on the information technology development and maintenance group. Whilst in any service setting it is beneficial to " own the customer" (Normann, 2001), it is clear that competition for time in front of the customer from different parts of the same organisation is unlikely to benefit the organisation or be a pleasant experience for the customer. Often there is an emphasis on the parts of the organisation which directly impact (Voss & Zomerdijk, 2007; Garcia, 2007) the customer (marketing)-sometimes referred to as the moment of truth, over those responsible for back-end delivery (operations)-information technology referred to as geeks. How understanding of requirements and comprehension of capabilities are exchanged within any organisation, and in a service setting, shown to the customer is clearly of great importance. Who within the organisation gets access and control of customer requirements and their expression will have a major impact on the ability of the organisation to deliver. This can be particularly challenging as it is often the case in complex service delivery that it is hard to distinguish between what the customer asks for, what the customer wants and what customer needs. How the understanding which resolves these issues is arrived at and transmitted will have major consequences for both the delivery of any service and the likely customer satisfaction. The long-term nature of service relationships requires an ability to look back at why a service design was believed to be capable of meeting a customer need and correct the design when it is no longer does so. To meet this goal both the

customer and the provider must have some degree of surety that the data underpinning the service design is actually valid. This requires that the data be explicit and in a form where all of the parties can agree that it represents valid comprehension.

Most service contracts enacted between a supplier and a customer have terms that express both what service is to be provided (function), how often it should work (availability) and with what speed (performance). These fundamental aspects of the service are expressed as service level agreements (SLA) and consequently guide the future relationship between the customer and the supplier. From both parties perspectives these SLA are essentially predictions. From the perspective of the customer they predict what the business needs of the customer are from the performance of the service. From the perspective of the supplier the SLA is a prediction of a viable economic capability, that is that the supplier can meet the SLA whilst still making a reasonable economic return (Taylor & Tofts, 2003). The question is to what extent does either party understand their reliance on the correctness of these predictions, or how they exchange information to establish their validity.

Given that the dominant requirement in the service setting is the ability to predict it is important to distinguish between predictive models and descriptive ones. Often the view is taken that all that is necessary is to measure something. Indeed some organisations have taken this to its logical conclusion and deem it to be the dominant requirement in any setting. Unfortunately the possession of detailed measurements of a system does not tell you how to correct that system if the measurements are not the ones you want. This is true both the business Systems themselves and the techno-social systems which support them. If we do not understand a service, then this is of little import when the service is working well, but will be a major impediment when the service is working badly! Furthermore, basing our approach to innovation within services on the measurements we have made upon them is likely to succeed more by chance and elasticity in the delivery system than by intent. Sadly on the modelling side there is often a temptation to arrive at the analysis that the systems are either chaotic or complex. Given the sensitivity to initial conditions of chaotic systems the next natural step is to conclude that such systems admit no useful prediction-interestingly whilst the motion of planets and the moon are chaotic and drive the tides it is perfectly possible to obtain cheap software which will tell you the height of the tides for all the world's ports for the next 200 years. Equally, the observation that the system is complex, usually when represented in its full detail, leads to the adoption of descriptive approaches since the complexity cannot (or is that should not) be overcome. However, it is the case that for most systems that are only a limited number of design choices available and exploiting these limitations within the comprehension problem can greatly simplifies it. Unfortunately since the analysed problem is no longer complex that has no impact upon the complex systems view.

In any complex service system the point at which the boundaries between human performance of tasks and automation of tasks lies has an enormous impact on the performance of the system. Inevitably the information technology centric view of service delivery attempts to automate everything, this aligns with the Ford-ist of view that a human is simply filling the space of the machine that you haven't bought yet. Unfortunately (Cross, 2002) in an increasingly automated system the consequence of a failure (exception) becomes increasingly costly. In an absence of understanding of the relative cost and performance consequences, particularly under exceptions, from the automated versus the human solution there is an inevitable attempt to over automate, as this *must* reduce cost. This may be that from a managerial perspective information technology is either perceived to be easier to understand, or can be ignored, than leaving humans in the system.

In systems that have multiple parties, either external or internal, which is inevitably the case for service systems (indeed many authors believe that the **co-creation of value is the distinguishing feature of services**) we inevitably desire to give an account of the systems co-evolution. Taking the information technology sector as an example comprehension work is often not undertaken because the speed of evolution within a single sector is so great that it is felt

impossible to understand. Coping with the pace of change of multiple interacting sectors is a challenge so great that it simply should not be undertaken.

Finally, the economic provision of services works! There is a vast and growing IT services sector which is extremely profitable, so where is the problem? This is comparable with the software systems view of design analysis, at its current extreme which is to develop and to test as rapidly as possible and largely ignore any issues of comprehension. Many within the IT industry are happy that this delivers economically viable products. It is unsurprising that such views transfer to the provision of service, given the level of IT presence in all aspects of service delivery. Within a service provision negotiation those comprehension activities that are undertaken will be dictated by the owners of the negotiation. Typically a sales and marketing organisation on the part of the provider and a purchasing organisation on the part of the customer. The presence of comprehension specialists in such negotiations is often seen by both of these parties as simply an impediment to their success. The salesman wants to make his sale, and does not want his negotiating strategy impaired. The procurer wants to make his savings and does not wish to reveal the full potential value to the supplier. Is there any value in a deeper understanding of the service provision to either of these parties?

Can we already solve the comprehension problem?

Are the problems with comprehending service systems symptomatic of a failure of current comprehension methods or do they indicate that the current participants are simply unaware of what is available? More fundamentally is it possible to have a genuine science of services without an underpinning theory? As an example, many people in Information Systems would believe that the theoretical aspects of their subject are almost entirely if not completely irrelevant.

Before suggesting that current methodologies for comprehension have failed one might ask what would constitute evidence for such a failure? Should the lack of use of particular techniques condemn them as failed? Does the fact that all of the stakeholders impacted by a service comprehension exercise may not understand the technique or even be willing to believe its results constitute failure? Should the fact that a particular methodology is not currently regarded as warranting research effort mean that it does not provide a viable comprehension technique? Should the fact that a methodology may only be visualised numerically mean that it cannot provide the basis for understanding service systems?

Indeed what actually constitutes current technique? We could regard something as a viable current method if certain proportion of stakeholders must understand it in order to make valid decisions in respect of the service system. Unfortunately this approach would be stymied by the various favoured techniques of the many disciplines which impact upon the service world, unsurprisingly each would choose their own. It is also unclear how widely comprehension techniques are actually taught. Whilst at that same time, there is almost no academic subject, which impacts upon services, which does not teach data comprehension techniques-commonly referred to statistics. Justifiable abstract comprehension techniques are the province of applied mathematics usually within the context of an engineering framework. In this setting it is clear that the validity of particular comprehension methodologies would largely be established by the existence of the community using those techniques to understand the problems of services. Even with IBM's strong call to arms within the services sciences management and engineering programme little has emerged with the above properties. Indeed, one might reasonably say that simulation techniques are the only approaches meeting the above requirements. However, the limitation of this approach has already been outlined.

Given the many activities which can inform the view of service systems it is unsurprising that each of them would claim that their dominant understanding approaches should be the ones applied. If one is seeking an abstract view then *a priori* the claims of systems engineering or operational research to "already have the answer" need careful consideration before they are

rejected. Equally, given the pressing nature of the economic problem perhaps we do have to accept that the only viable comprehension approach is through some form of in-silico emulation, and that our understanding, is inevitably bounded by forms of numerical experiment.

An area where current comprehension techniques are limited and not widespread is that of concurrency. Given the view that service systems inherently involve two or more players, one may conclude that the limitation of current abstract approaches derives from the presence of concurrency. In the setting of performance engineering, the presence of splitting a task into multiple parts each of which is worked on simultaneously (possibly with the various components interacting) and then conjoined in order to provide a solution, is an extremely hard problem[1]. Similar evidence comes from the problem of parallel programming, in comparison with sequential programming, where abstracting away from the underlying physical architecture almost invariably leads to a very poor (inefficient) solution. Consequently the exploitation of large-scale parallel systems tends to be via bespoke solutions. Even within the mathematical communities, who spend their time studying these problems there is almost no agreement as to what constitutes an approach which is both a valid representation and admits effective solution.

From the perspective of the applicability of current technique we have a similar problem as to that of attempting to generate the "master model". Much like the question should we have a collection of small models each of which enables us to understand one of the key interactions in the dynamics of the service delivery, is it overly ambitious to attempt to have a single comprehension technique which permits us to understand each of those key interactions. As an example it would be unusual if any service provision contract did not contain availability and delivery rate requirements, whilst both of these can be comprehended within the same technical framework the simplest approach is to treat them both separately. This observation may be one of the underpinnings of the difficulties of service comprehension, just as there are relatively few individuals (Monahan, et al., 2006; Spohrer, 2006) who understand the many disciplines which would underpin services sciences-the label mathematician does not usually indicate a generalist in abstract understanding, but more commonly a specialist in a particular approach. Consequently, whilst current comprehension techniques may well be up to the problem of dealing with the underlying issues in service systems, there may be almost no individuals who possess that range of skills.

Significant emphasis is placed within the emerging services sciences community on the co-creation aspects of services. Taking this as the dominant mindset which distinguishes the services world from the other economic realms the implication is that fundamental to understanding services is the ability to understand systems with more than one active participant. There are already two areas of mathematical development which take this view as fundamental, the theory of games and the theory of concurrent systems. Game theory (Conway, 1976) has been widely exploited to understand economic systems (indeed the easiest way for a mathematician to win the Nobel Prize) and has had some applications in services. The use of multiplayer games may well be particularly powerful in this setting. The theory of concurrent systems has largely been developed to deal with the problem of the correctness of communication systems, this has many embodiments ranging from actor semantics (Hewit et al ,1973) , through Petri nets (Petri, 1978) to process algebras (Milner, 1990) and beyond. However, these approaches have had limited traction even within the computer science community.

Taking a designed object approach to services can greatly simplify the service system comprehension problem. Traditionally the applied mathematical approach has been to accept the problem as stated and attempt to analyse it. Whilst this is entirely reasonable, nay essential, when dealing with the physical universe, this is actually unnecessary when dealing with human constructs. The principal of designing systems so that they are easy to maintain or so that they

[1] if it was easy the Danish queuing theorist Erlang would have solved the problem in the 1930s or Neuts in the 1970s (Neuts, 1995).

are easy to produce are well understood, however the idea that a designed system is constructed so that it is easy to comprehend is a radical change of mindset. Indeed the traditional separation (architect versus structural engineer for example) between the creative designer and the analyst positively mitigates against this outcome. An interesting question from the perspective of the customer is what would be their observation on a service system which was designed to simplify analysis? From the customer perspective one of the main demands is the reduction of risk, which can be regarded as an increase in predictability. Within a complex service offering there are likely to be parts where variation is irrelevant, equally there will be parts where departure from required performance will have major impacts upon the recipient of the service. Taking design for prediction as a principle this will enable better identification of which parts of the service drive cost and consequently a better identification of what the customer is paying for and why.

It is clear that there is a need for a collection of commercially (or government) relevant case studies of how service designs have been comprehended to demonstrate that they will meet their requirements. Unfortunately, in a commercial setting (even the government setting as they have commercial partners) the release of this data is unlikely to occur since it is central to any contract negotiation, even if the activities were not actually undertaken during the contract negotiation. There is an inevitable question of who would fund such initial studies and over what domains?

What are the research issues?

For a body of people with the relevant skills for service comprehension to emerge there needs to be a set of beneficial activities, both to them and to any funding organisation, for them to engage in. From the research funding bodies perspective these need to be concrete relevant questions which can be fitted within a well formulated scientific investigation. Equally, such investigations need concrete data against which they can be validated. Finally there needs to be clear linkage between the activities of service comprehension and those of service design to develop a fully fledged effective theory of service design.

One of the major issues from an academic perspective is the not unreasonable belief from many subject areas that this is already a solved problem-there is simply a lack of application of current knowledge. In this context generating motivating examples that illustrate why this is not the case, alternatively evidence why appropriate and focused educational materials needed to be generated, must be produced. Given the reward structures in academia if an activity is not recognized as research (within the academic's own subject) then it is hard to see why they should have any motivation to devote efforts to the area.

Consequences of service comprehension.

The major change in understanding services before we deploy them is that we would be able to move from a "poke and hope" approach to innovation and problem resolution, to an appropriate form of double loop learning. The recognition that the commercial contracts for provision of service are based on **predictions** which both parties rely on to as great a commercial extent as we rely on the predictions of structural engineers when we safely cross bridges, inevitably changes our approach to both contract negotiation and in-service governance.

The presence of predictive material means that we have the potential to undertake root cause analysis into service delivery issues. In particular this allows us to distinguish between when our assumptions were wrong, when the implementation was incorrect, when the models were inadequate or when the measurements were simply statistical artefacts. Such double loop approaches are fundamental to any continuous improvement technique. If we cannot distinguish between improvements derived from luck, and improvements which derive from judgement then we are unlikely to be able to deliver improvements with any certainty.

Most complex service systems involve multiple parties providing a service to at least one, but

potentially more than one customer. In such settings the governance of the relationships between the parties is inevitably difficult. Clear comprehension of what is required from each of the parties and what assumptions they have to make, so that the service can be successfully delivered changes the nature of their interaction. In particular we will understand what valid measurements are and can move to a predictive approach to service quality, rather than reactive one.

Acknowledgements.

Based on a discussion initiated by Chris Tofts at PICMET workshop on SSME 2007, with: Daniel Berg, Rensselaer Polytechnic Institute berga@rpi.edu, Antonie Jetter, Portland State University, jettera@cecs.pdx.edu, Masanori Namba, Ritsumeikan, Asia Pacific, nambama@apu.ac.jp, Matthew Realff, NSF, mrealff@nsf.gov and using the contemporaneous notes taken by Tad Drozdowski, Portland State University ETM Student.

All remaining errors are solely the property of the authors.

References

Behara, R.S. and R.B. Chase (1993). "Service Quality Deployment: Quality Service By Design." In R.K. Sarin (ed.), Perspectives in Operations Management - Essays in Honor of Elwood S. Buffa. Boston: Kluwer Academic Publishers, pp. 87-100.

Alan Bundy, personal communication, 1988.

John H. Conway, On Numbers and Games (London Mathematical Society Monographs), 1976

Michaal Cross, Kicking the habit, http://www.guardian.co.uk/technology/2002/oct/31/internet.epublic, Guardian Newspaper, 31st Oct 2002

Fausto Garcia: Exploratory Study of Management of New Service Development Process in Italian Luxury Hotels, http://www.poms.org/conferences/cso2007/talks/45.pdf, POMS/CSO, 2007.

Carl Hewitt, Peter Bishop & Richard Steiger, A Universal Modular ACTOR Formalism for Artificial Intelligence, http://dli.iiit.ac.in/ijcai/IJCAI-73/PDF/027B.pdf, IJCAI, 1973

Raj Jain, The Art of Computer Systems Performance Analysis, Wiley, 1991.

R. Milner, Communication and Concurrency, Prentice Hall, 1990.

Monahan, Brian; Pym, David; Taylor, Richard; Tofts, Chris; Yearworth, Mike, Engineering the Polymath, Innovation in Systems and Services Education, http://www.hpl.hp.com/techreports/2006/HPL-2006-100.html, HPL technical report, 2006.

Neuts, Marcel F. Algorithmic Probability, Chapman & Hall/CRC, 1995

Richard Normann, Reframing Business, John Wiley & Sons London, 2001.

Carl Petri, Concurrency as a Basis of Systems Thinking. St. Augustin: Gesellschaft für Mathematik und Datenverarbeitung Bonn, Interner Bericht ISF-78--06, Sep, 1978, Second Edition:, Proc. from 5th Scandinavian Logic Symposium, Jan., 1979, Aalborg / Jensen, F.V.; Mayoh, B.H.; Moller, K.K. (eds.) --- Aalborg: Universitetsforlag, 1979, Pages: 143-162

Scott Sampson: (Why We Need) An Operations Service Paradigm, POMS/CSO,

http://www.poms.org/conferences/cso2007/talks/02.pdf, 2007

Jim Spohrer, Education for Service Innovation National Academy of Science, Washington, DC, http://www.almaden.ibm.com/asr/SSME/esi/slides/Jim-Spohrer-NAS-Education-for-Service-Innovation-20060418.pdf, April 18, 2006,

Jim Spohrer, Service Science, POMS/CSO, http://www.poms.org/conferences/cso2007/talks/37.pdf, 2007

Richard Taylor & Chris Tofts, Minding the abstraction gap, POMS/CSO, http://www.poms.org/conferences/cso2007/talks/27.pdf, 2007

Richard Taylor & Chris Tofts, Modelling, myth vs reality, map vs territory, http://www.hpl.hp.com/techreports/2003/HPL-2003-246.pdf, HPL technical report, 2003.

Richard Taylor & Chris Tofts, Death by a thousand SLAs: a short study of commercial suicide pacts http://www.hpl.hp.com/techreports/2005/HPL-2005-11.html, HPL technical report, 2005.

Chris Voss and Juliana H. Mikkola, Thinking about service design, http://www.poms.org/conferences/cso2007/talks/49.pdf, POMS/CSO, 2007

Chris Voss and Leonieke Zomerdijk, Perspectives on Service Innovation and Design, http://www.poms.org/conferences/cso2007/talks/03.pdf, POMS/CSO, 2007.