Model Based Services Discovery and Management Industrial Research and Practice Paper

Richard Taylor and Chris Tofts
HP Laboratories, Bristol
HPL-2008-19
February 29, 2008*

Outsourcing, Services, Modelling, Scenarios, Risk

Service systems fail all too frequently. `Overdue, over budget and disappointing' are the words frequently used by organisations to describe their experience in the development and commissioning of complex information systems enabled services. More considered analyses question anticipated productivity gains, and in the longer term, a failure of service provision to track the changing requirements of the organisation.

As a major supplier of IT and IT-enabled services, Hewlett-Packard has invested heavily in developing and understanding of the reasons that services fail to delight, as well as developing technologies and management processes that mitigate against failure. This paper describes a (predictive) model based approach to service-systems analysis that aids in understanding the goals, the specifications and dynamics of a service system. Our contribution is a model based service discovery process and technology that can be used to dramatically improve inter-stakeholder communications, provide a design and management infrastructure that is robust to the inevitable changes that affect any commissioning organisation, and lay the grounds for more sophisticated cost-benefit analyses than are currently commonly used.

We draw on a number of large scale (multi-billion dollar) service projects to illustrate the application and benefits of this approach to service discovery and management.

Internal Accession Date Only

Approved for External Publication

Submitted to Portland International Center for Management of Engineering and Technology, PICMET 2008, South Africa

Model Based Services Discovery and Management
Industrial Research and Practice Paper

Richard Taylor & Chris Tofts
Systems and Services Research Centre
HP Laboratories Bristol,
{richard.taylor, chris.tofts}@hp.com

Abstract

Service systems fail all too frequently. 'Overdue, over budget and disappointing' are the words frequently used by organisations to describe their experience in the development and commissioning of complex information systems enabled services. More considered analyses question anticipated productivity gains, and in the longer term, a failure of service provision to track the changing requirements of the organisation.

As a major supplier of IT and IT-enabled services, Hewlett-Packard has invested heavily in developing and understanding of the reasons that services fail to delight, as well as developing technologies and management processes that mitigate against failure. This paper describes a (predictive) model based approach to service-systems analysis that aids in understanding the goals, the specifications and dynamics of a service system. Our contribution is a model based service discovery process and technology that can be used to dramatically improve inter-stakeholder communications, provide a design and management infrastructure that is robust to the inevitable changes that affect any commissioning organisation, and lay the grounds for more sophisticated cost-benefit analyses than are currently commonly used.

We draw on a number of large scale (multi-billion dollar) service projects to illustrate the application and benefits of this approach to service discovery and management.

1 Introduction

The many confusions between measurement and control in business are well illustrated by this (extended) quote from Michael Hammer[3] reflecting on his observations from a 'major electronics company'

> *The chaotic state of contemporary measurement was impressed upon me when I attended a senior executive meeting of a major electronics company, at which the company's leaders were carefully reviewing their dozen or so key performance measures.*
>
> *The executives meticulously examined a list of measures that was notable for its breadth: customer satisfaction, sales closure ratio, market share, order fulfillment time, employee satisfaction, working capital, service cost per customer, customer retention, new product break-even time, revenue per employee, and return on equity. Some of these numbers described overall company objectives (return on equity and market share), some were operational metrics (service cost per customer and order fulfillment time), some were miscellaneous items (employee satisfaction and customer retention).*
>
> *But what was most enlightening about the meeting and the list of measures was that the executives around the conference table had no idea what could be done to improve any of these numbers. If the numbers were good they would smile. If the numbers were bad*

they would click their tongues and make a careful note that something would definitely have to be done to improve that measure by the next executive meeting. Then they would move on to the next item.

The measurement system did not connect the numbers to each other in a meaningful way or provide executives with any guidance as to how to improve them.

Michael Hammer - Why Leaders Should Reconsider Their Measurement Systems, Leader to Leader, 24, 2002.

The catalogue of large scale IT outsourcing contract failures is a long and growing one. Recently it was estimated that the United Kingdom alone has wasted $18 Billion [6] on failed IT outsourcing projects. Alternative estimates suggest that 8% of government IT systems [5] in the UK are 'not fit for purpose'.

These outsourced service-systems place tremendous challenges on both their procurer[4] and their suppliers. These projects share a number of characteristics in that they:

- typically run for more than five years;
- are high value contracts ($400 million to $4 billion);
- will bebusiness critical – these IT services are central to customer success;
- have a large numbers of users, typically more than 5000 but in some cases as many as 200,000;
- are served by complex and large scale infrastructure investments;
- contain a combination of many (more than 50) lesser, but clearly identifiable services;
- will be managed by large teams of staff (in excess of 500).

Usually the main effort in specifying and negotiating the contract to supply these outsourced systems is concentrated on the system functionality, details of the infrastructure (hardware and software) and basic volumetrics. In our experience[9], even in such significant programmes, very little, if any time is spent on developing and demonstrating the value of the service components and their mode and dynamics of delivery to the business. Even less time is spent considering how assumptions about those behaviours can be validated and then tracked as the system (inevitably) evolves.

This becomes very obvious when the *Service Level Agreements* or SLAs that will be used to police and manage supplier-customer negotiations are examined. SLAs are the cornerstone of any outsourcing agreement and should capture two important properties:

1. the level of performance that is *required* by the customer to be succesful in its market and executing its forward business strategy;
2. an economically sustainable delivery point for the supplier, so that it can maintain its business (and means of tracking the inevitable changes to that business).

Unfortunately it is usually the case that neither of these properties are established in practice. A commonly used description 'best value' (which would presumably capture one and two above, is commonly simplified to 'lowest cost', sacrificing business value.

The first will fail because no business wants to expose the true value of a suppliers offerings for fear that the supplier will be put in a position where they can dominate the relationship. In the

absence of the establishment of a reasonable value on the part of the business it all of the suppliers simply bid as low as they can go - and hope to catch up on 'change'. Even more problematic is how customers distinguish suppliers on quality? What evidence can or should suppliers provide that they can achieve the delivery levels they claim and what are the consequences if they fail. Whilst there is a belief that penalty clauses for delivery failure can have the desired effect, there is almost no evidence of them succeeding [11, 9]

2 From SMART to SPART and beyond

> *What gets measured gets managed!*
> Traditional management proverb.

There is a well known standard approach to defining SLAs [2] which is usually refered to by its acronym SMART;

- Specific;
- Measurable;
- Agreed to (or Achievable);
- Realistic;
- Timely.

While these may be appropriate as a frame of reference for preparing performance systems for humans, this approach is limited (Figure 1) when dealing with designed systems[1].

In order to refine, and specialise the SMART approach to SLA definition and policing, we have taken a *control theory perspective*[8] of the business. This requires that SLAs should be integral to specifying the parameter space of the business that we wish to manage (Figure 2) - and at all levels within that business system, from objectives to infrastructure. This view of the business system has an important advantage over 'traditional' approaches to understanding, in that it makes it possible to *predict* the consequences of change within the model, and then respond to departures from that prediction. *Measurability* is therefore seen as being inadequate, what matters is *predictability*. There is little point in being able to measure something if one does not know how to control. Mere possession of the measurement without any understanding as to how it arose will at best lead to confusion and at worst provide a false sense of controllability.

If we begin from the premise that we need to predict outcomes, this implies that three questions must be asked;

- what do we need to measure?
- how often do we need to measure it?
- and with what accuracy do we need to measure it?

A supplementary question can be asked (easier in 'greenfield' operations than existing systems) is *how can the system be designed so that appropriate measurements, made at appropriate times and with an appropriate accuracy be used to control it?*

[1] Despite the best efforts of the education, and now genetic engineering, systems it is hard to argue that humans are designed.

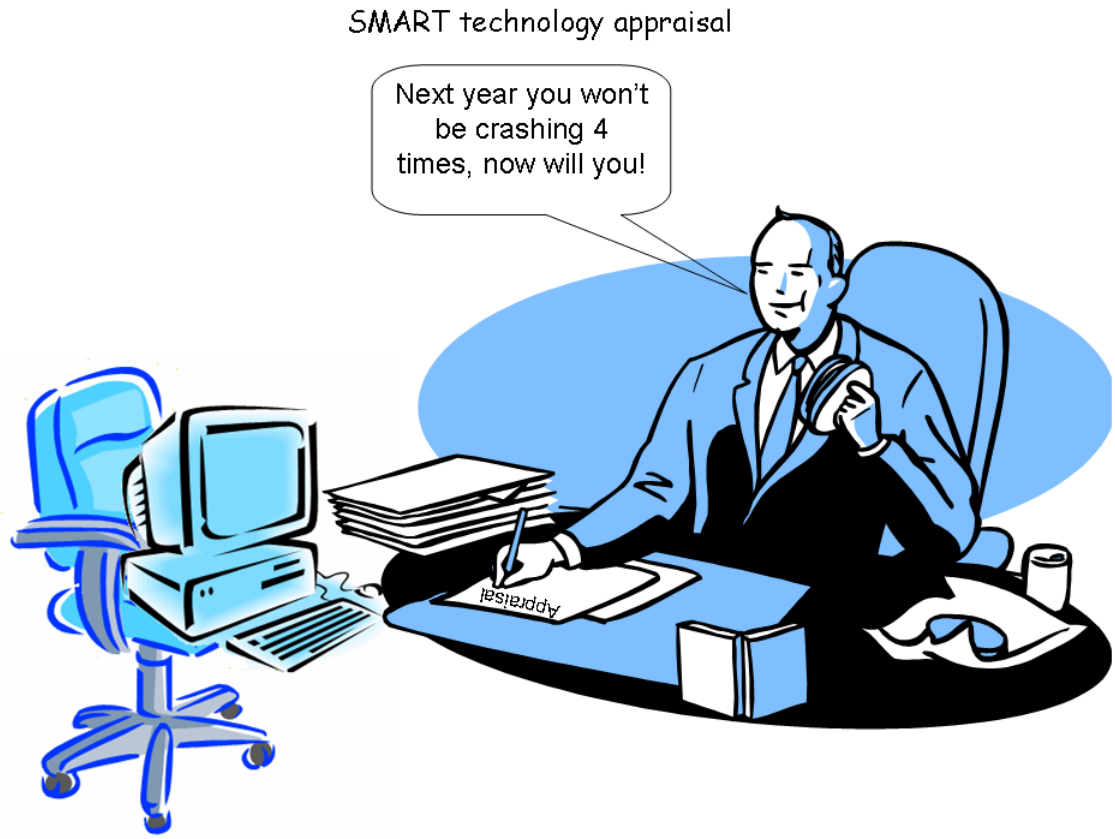

Figure 1: The use of SMART for technology.

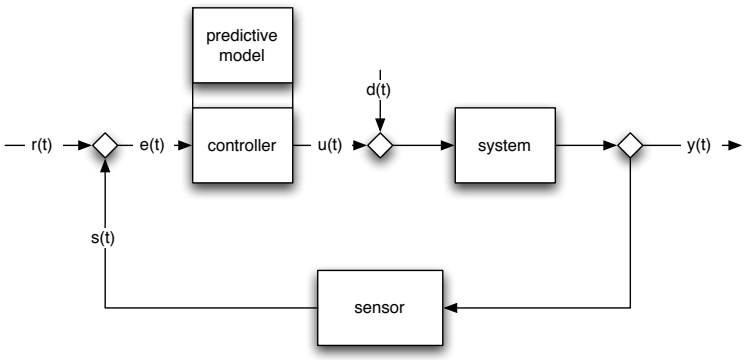

Figure 2: The control loop approach to systems management.

Most (but no all) large service systems have as part of their specifications, comprehensive lists of measurements that must be made of key performance indicators (KPIs) in order to demonstrate compliance with the SLAs that will be used to police the service. In our experience, none of these acknowledge the impact, either financial or technical of the measurement requirements. For example, we have seen instances of SLAs on inventory which would require the service provider to sample every single capital item monthly - or SLAs on network performance, whose establishment would consume between 10 and 40% of the available communications bandwidth. Clearly whilst measurability was achievable in both of these instances it was unclear what value it was returning.

2.1 Developing predictive models

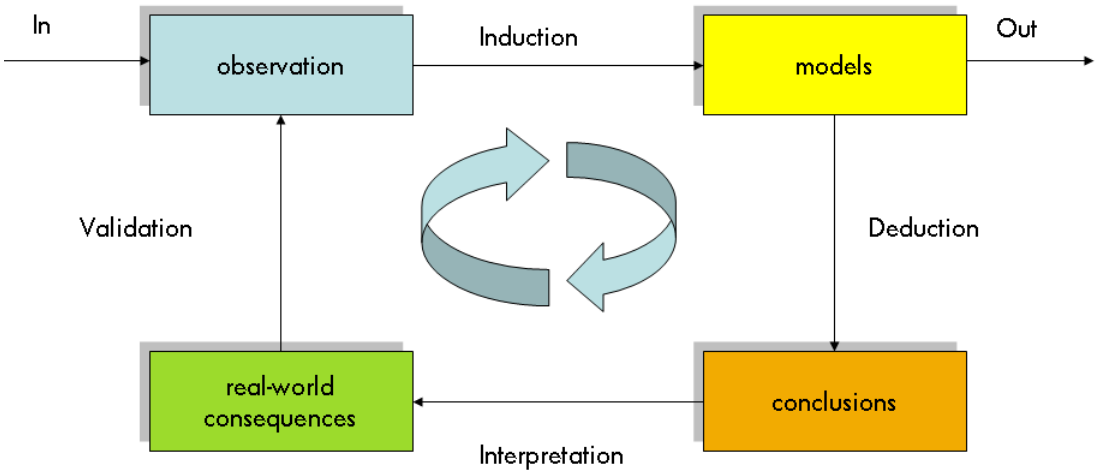

Figure 3: The development of models through the 'scientific process'.

The process is illustrated in Figure 4, this illustrates the progression from models of user behaviour, to models of allocation of users within a technical solution, to the requisite governance structures on that technical system derived from the two layers of models. Note, that as we have models of all levels of the sytem, we have a record as to why we need these levels of service for these particular users - consequently we have a starting point for rationally dealing with any subsequent changes either in user behaviour, user requirements or system capabilities.

3 Capturing the business dynamics - The RaSP methodology

> *...an organization needs to create a formal, structured, and quantified model of the enterprise – the kind that scientists and engineers use to describe physical systems. Such a model connects an organization's overarching goals with its controllable activities. Then, the organization needs to create a deliberate process for using measurement data to improve enterprise performance. This process must be structured and focused to use measurement information to identify the causes of inadequate performance and then do something about them.*

Michael Hammer - Why Leaders Should Reconsider Their Measurement Systems, Leader to Leader, 24, 2002.

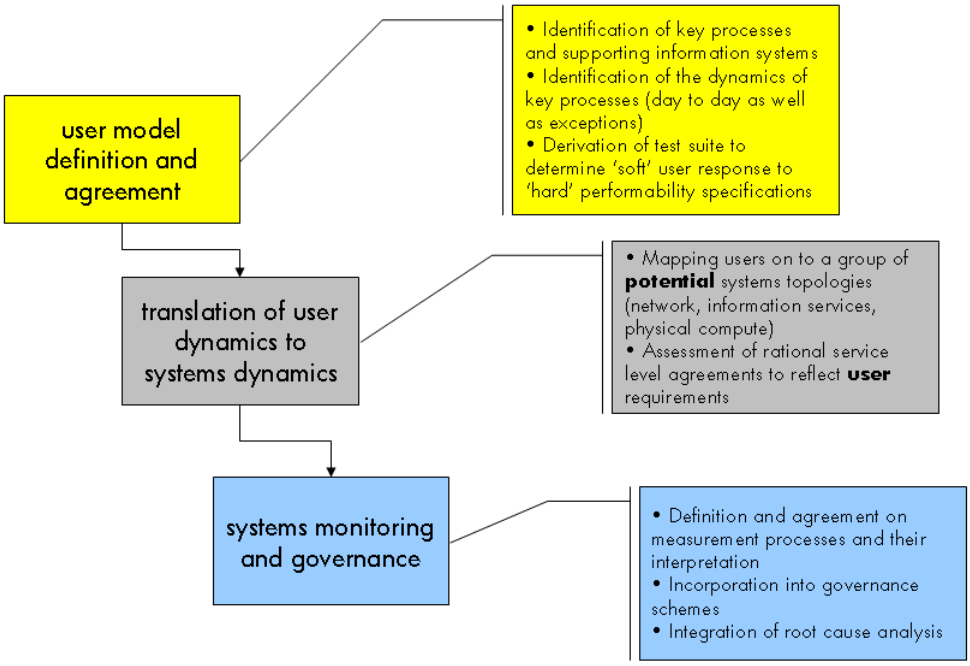

Figure 4: The overall process from users to governance structures.

Our development of the RaSP methodology has come about through work that has attempted to determine why some complex, but 'conventional' engineering programmes, based around established cross disciplinary teams succeed, whilst services programmes (even those of similar apparent complexity) fail. Our primary finding has been that a failure to communicate the consequences of (the) inevitable tradeoffs between design, implementation and management parameters between stakeholders[10] is a critical determiner. The best process engineering, human resources management and training will not fix problems that have been introduced through fundamental misunderstandings between the actors in the system.

Our work has concentrated on (but is not limited to) understanding and communicating the requirements of a complex information systems centric outsourcing bid process. The key concerns of such an activity are presented in Figure 5.

3.1 Implementing rapid scenario planning

Scenario planning is a well known approach to developing an understanding of the consequences of decisions - technical, social or business - on an organisation in an uncertain future. Most scenario planning processes involve facilitators who maintain conversations between stakeholders, prevent those conversations ossifying, and attempt to ensure that the space of possible decisions and their consequences is adequately covered. The results of such exercises are normally a combination of transcriptions of conversations that have occurred and informal (written) summaries of a few key outcomes and the expected conditions under which they might occur.

We have observed three primary weaknesses to this approach

1. the sequence of conversations that occur and their basis in evidence (hard numeric through to soft behavioural) is not captured in anything other than at best, transcripts and more

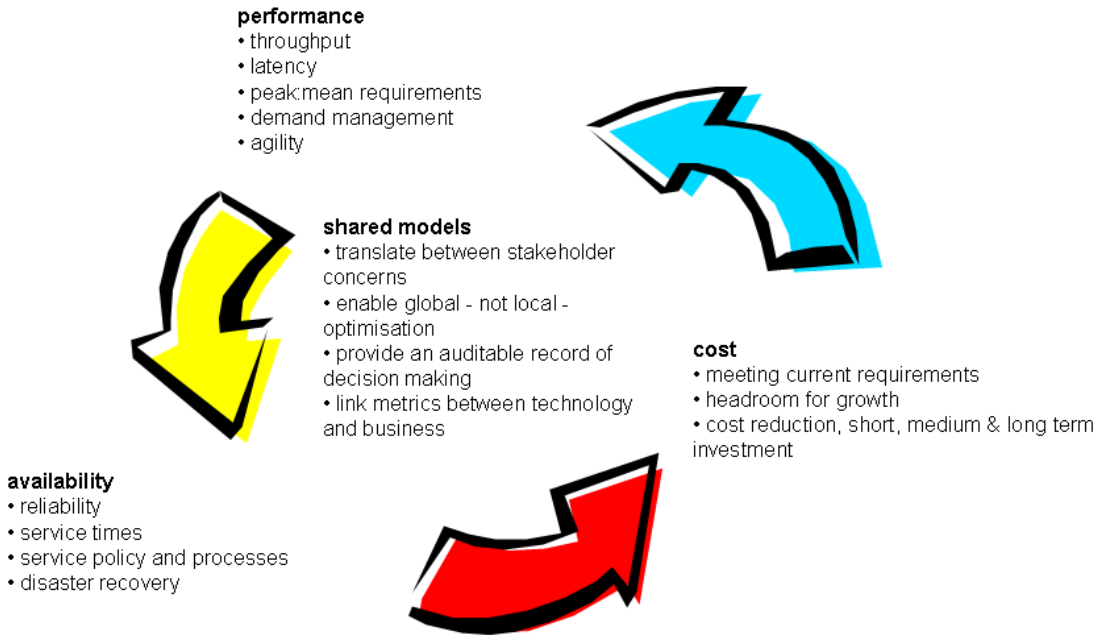

Figure 5: Key concerns for the RaSP process.

typically summaries of transcripts, flip charts and 'post-it' notes; in practice this leaves either no auditable decision trails or unwieldy and inaccessible records that are of little practical use;

2. the lack of an adequate evidence trail, along with the assumptions that have been used to drive the scenario outcomes are difficult to communicate to people outside the immediate process afterwards, and impossible 3-4 months later; given all of the important stakeholders will *not* have been present, this represents a significant barrier to communications, trust and in the presence of both success and failure, an ability to analyse what when right with the predictions and what went wrong;

3. the absence of an agreed means of following up the exercise and maintaining a rational and auditable conversation of refinement as new evidence comes to light further restricts the value of the exercise.

In order to address these shortcomings, we have approached the problem of predictions, audit trails and communications by centering the scenario planning process around sequences of formal, auditable models[7] that can be developed and their results demonstrated in real time, and which can be used to translate implications of decisions between different stakeholders who naturally have different[10] constituencies to satisfy. For example, in an e-Service, the marketing and sales, design, operations and financial teams will all have different and potentially conflicting interests when the dynamics of the underlying service are determined.

This approach leads to sequences of scenarios, backed up by predictive models that capture and document the conversations (Figure 6) between stakeholders (in the example used above these could include *our customers really care about the response times at this stage in the process*, *peak demands like that at the end of the month can't be met unless we re-engineer our procurement and*

support strategy and that will cost... because... to *the customer lifetime value is driven by these key factors... what can we do to optimise them?*

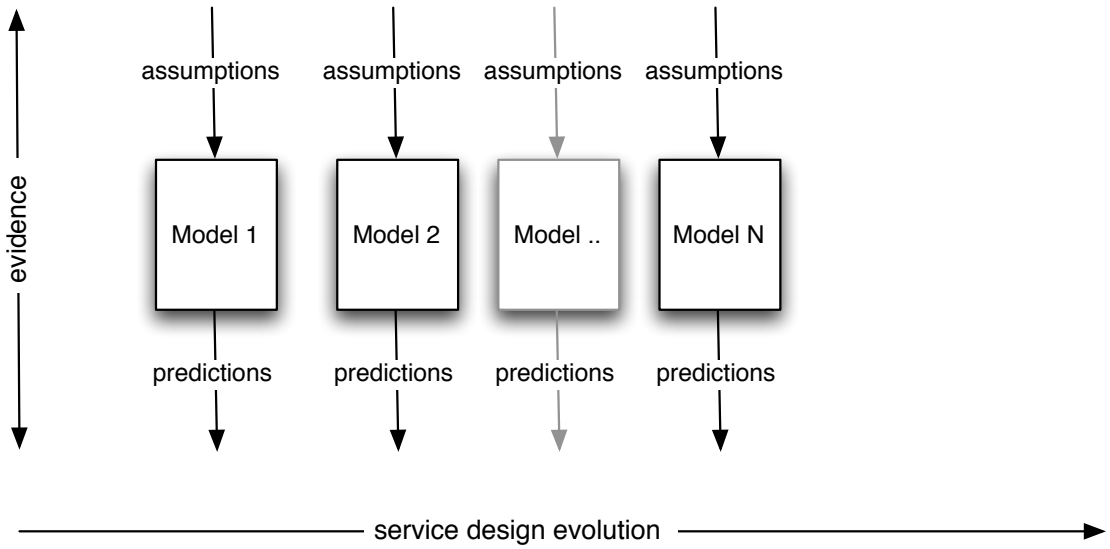

Figure 6: The development of sequences of models, with documented assumptions and outcomes is used to support the design, review and delivery processes

As well as acting as an audit trail (reasons for decisions and not just the decisions are captured and evidence - however uncertain is presented), these models can then be taken outside of the seed scenario planning activity and shared amongst stakeholders. This means that other groups can test predictions and question evidence, as well as maintaining a common language for exchanging their expectations. From the point of view of forward planning, the models make explicit the measurement-management loop, and also provide a basis for determining the value of KPIs and hence the appropriate service level agreements and governance structures.

Our approach has largely exploited a simulation oriented modelling language [1] which has advantages of conciseness and a user dynamic that makes it possible to derive and refine models quickly, although the RaSP process does not rely on any specific modelling technology. Our experience [7] however is that large, unwieldy and over complex modelling technologies will fail in this context, either because they are too slow, because the participants do not understand the assumptions embedded within them or because they require very specialised skills for their analysis.

3.2 Applying RaSP

The RASP methodology proceeds through five stages:

1. identify key stakeholders from target customer - these may include
    - user community representative
    - developer (sometimes known as R&D)
    - operations
    - business analyst

- marketing
- financial controller

but remembering to keep the group small - twelve is a crowd, the lower[2] the number of key stakeholders, the better the focus.

2. Interview the key stakeholders, preferably 2-4 weeks ahead of primary workshop to identify concerns, typically these will emerge from discussions about:

   - cost (fixed and variable)
   - risk (implementation and operations)
   - performance
   - availability
   - disaster recovery
   - agility (performance, availability, functionality, platform)
   - technology migration
   - operations
   - security
   - environmental impact

3. Prioritise issues and create a set of boiler plate performance, availability and cost models that illustrate key issues for each stakeholder group; choose appropriate language that reflects the stakeholder experience of the proposed service.

4. Hold a one day RaSP workshop. Attended by representatives of the key stakeholder groups and facilitated by a team consisting of:

   - Chair (who will have managed the pre-workshop survey)
   - Modeller 1 and Rapporteur
   - Modeller 2

   An outline timetable for such a workshop is:

   - Introductions
   - Key issues and expected outcomes
   - Workshop sessions - typically no more than 30 minutes a piece and broken up with 15 minute breaks (coffee, cigarettes and impromptu peace talks)
   - Wrap-up: drawing primary conclusions about requirements including areas of doubt and uncertainty, fixed points and unknowns with appropriate ownership
   - Fix expectations for deliverable

   For each of the workshop sessions of 30-45 minutes follow the following process:

   - fix on one issue for the session, dont worry, coupling will occur
   - prepare a use case scenario for each stakeholder

[2] A la Einstein everything in that needs to be and nothing else - sadly an identification made through practise.

- begin with an illustrative model that demonstrates the key issue - for example, availability vs cost
- remain focused - remind participants that this is not the only discussion being held, the intention of this session is to explore one key concern from each stakeholders point of view
- conclude the session with a statement of what has been agreed, what remains to be decided, what data needs to be gathered in order to make the decision, and evidence for that decision generated within the workshop
- after 30 minutes, if agreement is not possible, the conflicting stakeholders need to be tasked to refine their case (it might be appropriate to organise an independent meeting iff the conflicting stakeholders can be isolated from the rest of the group)

5. Engage the stakeholders in a followup activity:

   (a) report 0
   - all stakeholders are recipients
   - presents a uniform framework for the service (or SLA set) that explicitly describes the relationships between stakeholder priorities and their interaction
   - lists points of agreement and design & operations priorities
   - list points of conflict and/or data dependencies, as well as agreed actions to resolve these

   (b) report 1..N
   - individual stakeholders are recipients
   - presents key decisions in stakeholder specific language
   - lists agreed actions

The most important result of the RaSP process is that *all* of the shareholders are in receipt of models, which they can 'run', and that illustrate the dynamics and the assumptions that support the points of agreement and discord between them.

3.3 Tool Support

The RaSP process places particular demands on the tools[3] the modellers use to support the process. The aim of the workshop is to evolve the models in the presence of the stakeholders to reflect their views of the dynamic requirements upon the system. The following principles seem to apply to tool support for this space:

- lends itself to an abstract view of the system;
- does not hide assumptions;
- allows rapid analysis of multiple versions of a model;
- allows multiple views of the data produced by model;
- is presented in a form which is widely comprehensible;
- is in a form which *ALL* stakeholders can take away and play with;

[3]There is is an experimental licence version of one of the tools we use DEMOS2k available from www.demos2k.org.

3.4 Experience in Use

We have deployed the RaSP process in outsourcing programmes with a total value of over $20Bn. The universal view is that the methodology has greatly clarified each organisation's goals (including contracts in which 'multi-sourcing' from different vendors has been a feature of the work) within the contract. These have invariably lead to the reformulation of key performance indicators, governance processes and associated service level agreements in order to better reflect the customers real needs.

Allowing all parties access to the models changes the dynamic form *this is the best solution because we say so* to **this is the best solution we can find and here's how we found it, if you can find something demonstrably better then great, let's go with it**. Turnaround within the process is fast – We have examples the time between initial engagement of the HP RaSP team, and the delivery of models and primary conclusions to the client, including the the workshop, has been less than a week.

4 Model Based Governance - Keeping it Running

As with puppies, predictive models and the history of their formation are not just for christmas. In fact they have continuous application throughout the lifetime of the service, for both customer-supplier governance and also for internal supplier review.

4.1 Customer-supplier governance

No service design, however good, survives first (let alone prolonged) contact with the customer. Assumptions about the service utilisation, 'hidden' exceptions, and unanticipated customer expectations amongst other factors force changes in the delivered service. Over the longer term, customer requirements change – the business changes and business dynamics evolve with the service.

These mismatches between the delivered and the required service need to be tracked and rational decisions made about what modifications to the service (or indeed the way the service is being applied within the business) should be made and which parties are responsible. In the absence of prediction, and a means of relating the difference between prediction and measurement to service structure and dynamics, this is difficult if not impossible.

Throughout our service delivery, measurements (derived from the control theoretic analysis of the service) are matched against the predictions made by the multiple models that have been constructed and refined in the design phase. As the measurements inevitably depart from predictions then either the delivery model needs to be changed, the models themselves need further refinement to match the reality on the ground, or, as can be the case, the departures are shown to be artifacts rather than significant departures.

This enables a massive improvement in the governance process. The implications of deviation from expectation can be more effectively explored than before, enabling all of the parties involved to establish responsibility for change, if change is actually necessary. Furthermore, as either the customer evolves, or new technologies and processes can be made available by the supplier, the models provide a foundation for exploring their implications to the delivered service. This all represents a sea change from the 'traditional' approach to problem resolution (as each party draws their lawyers) as well as evolutionary service planning and design.

4.2 Internal supplier governance

Good services design is a combination of effective customer analysis, due diligence and appropriate delivery. Services that fail (either to provide the customer with an effective solution, or the supplier with acceptable margins) will be torn apart in an effort to understand what has now worked within

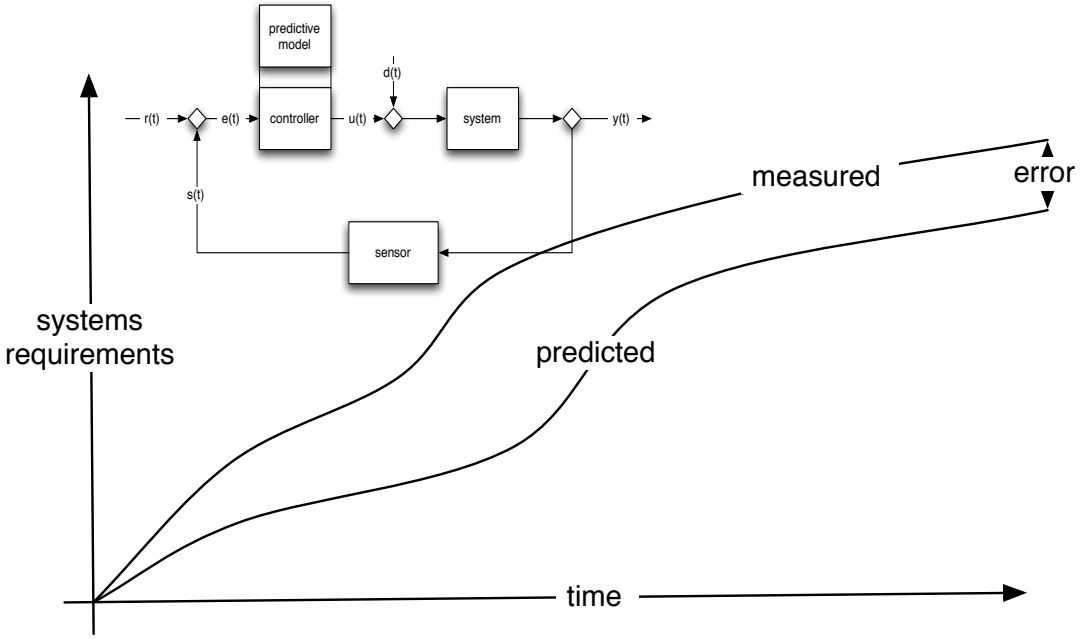

Figure 7: Governance process in the presence of prediction.

the analysis, design and delivery phases of the project. Less frequently (as they tend to attract less senior management attention), projects that succeed are pulled apart in order to develop 'best practices' that can be applied to new contracts.

In the absence of an auditable design trail that has captured assumptions and service structures, identification of both good and bad decisions is difficult. This is especially the case when the significance of those decisions needs to be understood. Human beings are remarkably good at identifying patterns and then assuming that correlation is cause – witchcraft has a long and (dis)honorable history, and not only in the services industry. Our model based approach to service design and delivery enables an effective evidence based approach to supplier review and continuous improvement than has been possible in the past. In part this is discipline (the evidence of decisions and their implications on design has had to be collected as part of the process), and in part it is the ability to review a project and explore what might have happened if other questions had been asked or other assumptions made.

5 Conclusions

There are many benefits to taking a predictive model based approach to the analysis, design and delivery of complex outsourced services, including

- it establishes a link between what the end user really cares about - be that monetary return on investment (commercial services) or social consequences (for public services), and the infrastructure (people, plant, information) that will be used to deliver that;
- reduced stakeholder dissonance – we have demonstrated and measured significant improvements in the inter-stakeholder communications processes, while differences between the mul-

tiple parties that have an interest in the service still remain, they are more clearly understood than was previously the case in conventionally managed projects;

- bounded risk-objective profiles – futures exploration with predictive models enables the integration and assessment of risk in the context of an organisations objectives and enables rational comparisons between apparently (dis-)similar service programmes;
- a basis for rational negotiation between parties – auditable, analytic models provide an excellent starting point for multiple parties to begin an exploration of delivery-expectation deviations. The rigor of the model based planning, albeit lightweight, provides an evidence base that is far more reliable than the participants hazy recollections of diligence and negotiation exercises that might have taken place as many as three years previous to the service introduction;
- increased 'fit for purpose' of solution – solutions designed and delivered through this process are demonstrably higher quality than historical approaches;
- purpose identified and documented – surprisingly (or not), the process of exploration exposes the service planning and requirements to a scrutiny that has often not occurred at the point at which a customer has decided to go ahead with the programme. Service objectives, value and risk are clarified, and we have observed contracts in which a no-go decision has been made subsequent to this planning process.
- measurement, monitoring and governance regimes that can be used by all of the stakeholders to assess value and manage change are an important outcome of the process;
- a basis for rational, repeatable decision making – many large services programmes involve considerable customisation of what were originally 'standard' offerings. Comparing and contrasting such programmes has been likened to comparing chalk with cheese – superficially similar until you bite into the detail. Our model based approach to specification and design enables appropriate comparisons to be made of structurally similar systems;
- a reduction in the need for design re-work – we have demonstrated significant improvements in the quality of design which has well known implications for the cost and stability of the delivered service;
- finally, this process returns control into the hands of those being served by the information system – something all to often forgotten, the 'poor bloody users' are given a far greater influence on the service that they will have to live with than has often been the case in the past.

References

[1] Demos2000, `www.demos2k.org`

[2] Peter Drucker, The Practice of Management, 1954.

[3] Micheal Hammer - Why Leaders Should Reconsider Their Measurement Systems, Leader to Leader, 24, 2002.`http://www.leadertoleader.org/knowledgecenter/journal.aspx?ArticleID=120`

[4] Aris Komporozos-Athanasiou, A grand opportunity in the IT services outsourcing market: Understanding the asymmetries of the client- provider relationship, HPL-2007-125 http://library.hp.com/techpubs/2007/HPL-2007-125.html

[5] Gov departments need better data systems, NAO says http://www.theregister.co.uk/2007/12/21/nao_data_systems_report/

[6] Outsourcing overruns cost UK taxpayers 9bn - http://www.theregister.co.uk/2007/12/27/uk_public_sector_outsourcing_cost_overruns/

[7] Richard Taylor and Chris Tofts, Modelling, myth vs reality, map vs territory http://library.hp.com/techpubs/2003/HPL-2003-246.html, 2003

[8] Taylor, Richard and Tofts, Chris. Business as a Control System: the essence of an intelligent enterprise. HPL-2003-247 www.hpl.hp.com/techreports/2003/HPL-2003-247.pdf

[9] Richard Taylor and Chris Tofts, Death by a thousand SLAs: a short study of commercial suicide pacts http://library.hp.com/techpubs/2005/HPL-2005-11.html, 2006

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

[11] Accenture escapes £1bn penalty for NHS walk-out http://www.channelregister.co.uk/2006/09/29/accenture_nhs_penalty/