White Paper

Services Research in Australia: A Roadmap

Aditya Ghose
Decision Systems Lab
School of Computer Science and Software Engineering
University of Wollongong
aditya@uow.edu.au

Services Research in Australia: A Roadmap

Aditya Ghose
Decision Systems Lab
School of Computer Science and Software Engineering
University of Wollongong
aditya@uow.edu.au

October 11, 2012

INTRODUCTION

Service science addresses one of the grand challenges of our times. It is a response to two inexorable historical trends. First, organizations are being obliged to *do more with less*. Organizational productivity has, in general, trended upwards as human society has evolved, aided both by *technological* and *methodological innovation*. Second, driven by increasing competition, organizations are being obliged to *improve the quality of engagement* with their stakeholders. Coupled with these historical trends are some *discontinuities* – unique to our times. These include the climate change challenge, as well as the enormous growth in (the global) population and an exponential growth in the demand for resources. In essence, these phenomena impose on us two significant imperatives: the *efficiency imperative* and the *quality imperative*. The former has been addressed to varying degrees by disciplines such as operations research, industrial engineering and computer science. The latter has been the focus of disciplines such as management and marketing. As a consequence, these questions have been addressed in a piecemeal fashion. That has impeded progress, because these two imperatives are intimately inter-related, and we cannot address one while ignoring the other.

Service science is probably our first attempt to understand enterprise functionality in an integrated, holistic fashion. Service science seeks to understand the way an enterprise conducts its business (or performs its functions), by using a service as the unit of analysis. The service in question might be delivered via the IT infrastructure – these services are of interest to the service-oriented computing community. Alternatively, these services might be human-mediated; these are typically the focus of interest of the services management and marketing community.

Traditionally, there has been little attention paid to the development of a comprehensive body of knowledge to support the design and delivery of enterprise functionality (the disciplines of management and operations research consider some, but not all, of these issues). Service science addresses that gap, by synthesizing results from computing, information systems, management and operations research (amongst several others). It includes in its ambit of inquiry questions such as the following: How do we model services? What components of enterprise functionality are best packaged as a service? How can we obtain a servitized view of the enterprise? How do

we identify candidates for outsourcing? What are best-of-breed methodologies to support the development of service designs? How do we improve service designs? How do we manage the complete life-cycle of a service? How do we ensure that the services offered by an enterprise are aligned with its strategic objectives? How do we ensure that services comply with applicable laws and regulations? How do we optimize service delivery? These are just a representative subset of a much larger repertoire of questions that this community addresses.

Internationally, there is a growing recognition of the value that such thinking might deliver. Societies and organizations similar to the Service Science Society of Australia exist in a variety of geographies, such as the Services Research and Innovation Institute (SRII), which is an international body but with a charter similar to that of the Service Science Society of Australia, the Service Science Society of Taiwan, and so on. A number of research institutes, such as the European Services Research Institute and the Karlsruhe Services Research Institute, are engaged in R&D in this space.

Australia can boast of an active and eminently successful services research community - spread over the university sector, CSIRO, the Smart Services Cooperative Research Centre and NICTA. Much of the research output of this sector is world-class. In understanding how and where to direct research investment and effort, it is important to ask: *How well does services research in Australia serve the needs of Australian society?* The answer is not as resoundingly positive as one would like. To varying degrees, Australia's services research community tends to follow a research agenda set elsewhere, often by larger research communities in other geographies.

This, then, begs the following question: *Is there a uniquely Australian research agenda?* The often-ignored answer is yes. Australia is situated in a unique economic context, with a unique set of opportunities and challenges. Australia's enormously successful mining and resources sector, its vibrant agricultural sector, its large and growing export-oriented higher education sector and its ageing population are unique, and require bespoke solutions from the services research sector.

We then need to address the question: *What kind of services research will have the most impact on the national economy and well-being?* The answers will span the spectrum from services management and marketing through to service-oriented computing.

Finally we need to ask: *What is necessary to enable/execute this research agenda?* There is a need for a clear plan of action to realize the value proposition that the services research community in Australia could deliver.

In summary, this discussion paper seeks to address these questions:

- *How well does services research in Australia serve the needs of Australian society?*

- *Is there a uniquely Australian research agenda?*
- *What kind of services research will have the most impact on the national economy and well-being?*
- *What is necessary to enable/execute this research agenda? How can the services research community deliver value?*

There is some urgency associated with these matters. Australia's mining and resources boom might have run its course. The manufacturing sector is in decline. The export-oriented higher education sector faces serious challenges. There are continuing questions regarding the efficiency of government service delivery and the cost of government. The healthcare needs of Australia ageing population will remain a major (but well-recognized) challenge.

The remainder of the report addresses these questions in some details. Much of the discussion involves highlighting what the service science community *can deliver*, but also what it *could deliver*. This is therefore a mix of a capability statement and a discussion of open questions.

SERVICE-CENTRIC GOVERNMENT

Reducing the cost of government has become a major concern everywhere but is especially true in Australia. There is recognition within government, at various levels, for the need for methodological innovation, but attempts at achieving this have been piecemeal, and have suffered from a lack of reference to a theoretical framework that would underpin such efforts.

An example is the attempt to reduce redundancy in internal functionality by centralizing replicated functions that can be found in a wide variety of government agencies (such as human resource management, accounting, logistics etc.) within "shared service" entities (such as the Shared Services Agency of the Government of Queensland, or the NSW Services First initiative). These are laudable efforts, but do not go far enough. A far larger set of functions could be shared if the unit of analysis was made finer-grained. The sharing of finer-grained functionality comes with associated coordination costs, but with an underpinning conceptual framework, of the kind service science could offer, trade-offs could be analyzed over the entire space of design alternatives.

A key requirement for analysis of this kind is the ability to model (or document) government services (both internal and citizen-facing) in a sufficiently expressive, but standardized format. The service science community is able to offer such service modeling standards. Coupled with elements of enterprise modeling and enterprise architecture techniques, these could form the basis for *whole-of-government service architectures*. This would support the identification of intersecting functionalities and capabilities within disparate parts of government. It would also support the strategic alignment of these capabilities. Several parts of government involve unique combinations process-driven and rule-driven execution architectures. Consider for instance tax

or social security entitlement processing. Many of the relevant services in these domains involve the execution of well-understand business processes, but a substantial number also involve highly flexible, rule-driven (and, in particular, legislation-driven) processing. Understanding and leveraging such insights in service design and execution poses interesting challenges.

Such government service architectures would also provide a critically-needed basis for change management. When changes need to be made to particular services, a range of analytical tools developed by the service science community could be brought to bear on identifying the impact of the change across the whole of government. This in turn would provide the basis for informed decisions on whether and how to implement change.

These techniques would also support the principled design of citizen-facing services. The service science toolkit offers answers to questions about what the right set of services to expose to citizens might be, what the appropriate service bundles might be, how service touchpoints might be designed and how value co-creation might be best supported in citizen-facing services. Ultimately, these techniques also provide the basis for the *normative design* of government service architectures, encompassing both internal capabilities and citizen-facing services.

Any framework that supports design must also support the maintenance of designs in the face of change. A key trigger for design maintenance is the need for service improvement, which in turn relies on machinery for monitoring and assessing the quality of the functionality delivered. The service science toolkit includes approaches such as sentiment mining of social networks to understand citizen responses to the quality of government services. The deployment of techniques for assessing quality of service can also add value in helping understand the effectiveness of government services and programs that indirectly impact citizens, and where citizen responses or sentiment does not form a viable basis for assessment. For instance, the effectiveness of a government program of investment in research in a particular domain is not easily measurable. Yet service science techniques can help us reason backwards over a network of causal connections from indirect manifestations of the impact of such programs to principled inferences about the effectiveness of such programs.

Distinct from *service design optimization* (as discussed above) is *service delivery optimization*. This involves answering questions about who delivers what service, when and how. These questions have traditionally been the preserve of disciplines such as operations research, industrial engineering and computing. Service science thinking adds value to these approaches by asking these questions, not in isolation, but in the context of the others discussed above. Thus, service delivery optimization must be addressed in the context of the design of government service architectures, as well as detailed citizen engagement models. Service delivery optimization in the enterprise context also throws up new challenges that the operations research, industrial engineering and computing communities have not yet considered, such as the *optimal*

service provisioning problem. Consider the following example: given a specific service, its detailed design, expected service demand and some service-level guarantees, what is the optimal team size necessary to deliver this service without violating the service-level guarantees? The answer to this question relies on machinery that performs *optimal task allocation* – given a service request, what is optimal allocation of the tasks involved in processing this request to individuals or roles?

ENTERPRISE SERVICE MANAGEMENT AND INNOVATION

Almost all of the issues discussed in the previous section apply as much to the government sector as to enterprises in the private sector. There are, in addition, challenges that private sector firms need to address (as with the others, these challenges also apply to government, but the case is more compelling for the private sector).

A key competitive challenge is *service innovation*. The need to constantly re-engineer and re-invent the business and the need to constantly seek new strategic positions to leverage or strategic resources to develop are part of age-old management wisdom. But for Australian firms facing increased competitive pressures, the need to innovate in service design and delivery is particularly acute. A number of firms have tried the so-called "open innovation model" that seeks to incentivize internal innovation using techniques such as enterprise crowdsourcing. The reported evidence suggests that these efforts have met with mixed success. Much more research is needed to understand and identify a richer range of models for supporting service innovation, including models for *systematizing* and *institutionalizing innovation*.

A major challenge for Australian firms is the management of outsourcing - and sometimes, offshoring – which is often driven by acute cost pressures. The most common form is the "outsourcing" of a firm's IT infrastructure and its maintenance to a cloud service provider. Service science thinking offers solutions to the problems - both technological and methodological (such as risk management) - involved in moving enterprise systems to the cloud. Sometimes the outsourcing of entire business processes is an option. New models of business process outsourcing, such as "people clouds" (where extra-organizational pools of human resources can be flexibly provisioned as needed) are making the space of decisions both richer and more complex. "Cloudbursting" – a cloud computing model where an organization accesses computational resources from the public cloud to manage potentially short bursts of transient demand – is now a possibility with "people clouds". This is a dramatically different version of the traditional business process outsourcing model, and deep insights will be required to understand how to best leverage potential efficiencies.

More generally, businesses find themselves having to manage mixes of IT- and human-mediated services of increasing complexity, and the in *open systems* settings (where system components change dynamically and unpredictably). Yet the business-IT divide remains as large as ever. IT

components of most modern businesses are managed using a set of principles totally distinct from those governing the human-mediated components. There is a growing body of work within service science that is taking a fresh look at the notion of socio-technical systems (the concept itself has been around for several decades, but the current emphasis is entirely new). A likely medium-term service science research outcome of value to Australian businesses will be new insights into the management of these unprecedented peer-to-peer compositions of humans and IT systems.

Another key business need is support for compliance management. With increasingly stringent legislative and regulatory frameworks on the one hand, and increasingly onerous costs of compliance, on the other, businesses need help with achieving compliant service designs, monitoring service execution and achieving design/execution repair in instances of non-compliance. These are questions that the service science community has considered in some detail (e.g., frameworks for business process compliance).

The computing dimension to service science offers a ready-to-deploy technological tool-set of value to Australian organizations. The key value propositions of service-oriented computing and SOA (service-oriented architectures) involve mitigating the effects of the business-IT divide, achieving flexible software architectures, ease of maintenance and so on are well-known. The existing and mature, research agenda on service-oriented computing will no doubt contribute to supporting the rapid and flexible deployment of IT systems in Australian firms in a manner that is better correlated with core business objectives.

SERVICING THE RESOURCES SECTOR

The resources sector in Australia has been one of the major drivers for economic growth in Australia in the recent past. Yet, the sector faces challenges that the services science community might be able to help mitigate.

There are possible future scenarios where the Australian mining sector might come under severe cost pressures. Some have recently argued that such a scenario has already eventuated. Leveraging service science innovations to obtain productivity gains and cost efficiencies would be key to any strategic response under these circumstances. Service design and delivery optimization techniques could be deployed to improve mining (or oil or natural gas) supply chain efficiency. Given that the resource sector has ridden a wave of high resource prices for a fairly long time, supply chain planning and optimization have been under-emphasized. Hence, even small efficiency initiatives, targeting “low-hanging fruit”, have the potential to make a significant cost impact.

Another potential response to the cost efficiency challenge is leveraging outsourcing. In the mining sector this is commonly described as “contract mining”. In this model, key mining

operations, particularly those involving capital equipment, are outsourced to subcontracting entities. For instance, the traditional operating model for an open-pit mine had the mine operator making major capital investments in heavy excavating machinery, such as shovels, draglines etc. In the contract mining model, excavation service providers would own such machinery, and would enter into contracts with the mine operator to operate and maintain these machines, with service-level agreements that would guarantee a certain level of throughput (e.g., in terms of tons of rock per hour), maximum acceptable levels of equipment downtime, conformance to occupational health and safety standards, and so on. Contract mining is also used in the operation of dump trucks in open-pit mines, with similarly specified contracts. It is clear that the scope of contract mining can be significantly extended to include, in effect, all mining operations (including those with high associated risk, such as drilling and blasting, that mine operators are often reluctant to outsource).

The key impediment to achieving better cost efficiencies via a more widespread deployment of contract mining has been the problem of *servitization* – the packaging of enterprise functionality in a (conceptual) "service wrapper". The notion of a service wrapper derives from service-oriented computing, and involves the "wrapping" of a legacy IT application in a layer of code that makes that application accessible via a standardized service description language, in turn enabling the use of that application within a service-oriented architecture, or invocation as a web service. One can also conceive of service wrapping for non-IT business services or for socio-technical services, where these functionalities are specified in a standardized format that permits service-oriented analysis, monitoring and outsourcing. The process of servitizing a given enterprise function provides a principled basis for specifying the contract that would govern the outsourcing of that function. Inadequately specified service descriptions lead to flawed service contracts – many outsourcing contracts, including some for contract mining, have suffered thus. Contract mining requires specialized monitoring, where again servitization plays a critical role.

As discussed above, there are possible economic scenarios in the not-too-distant future where the current runaway commercial success of the Australian mining sector might be somewhat attenuated, or where the sector's substantial contribution to the national economy might be somewhat diminished (which might be caused by a range of factors outside Australian control, including cyclical market downturns, altered economic circumstances in some of the principal overseas markets as well as altered geopolitical scenarios). None of these scenarios oblige us to assume that there will be a *global* downturn in mining investment – in other words, demand for mining innovation will persist in other parts of the globe (for instance in the mining sectors of large emerging economies which primarily cater to domestic requirements). In these scenarios, the Australian mining sector could continue to thrive as a mining *service provider*, even if its role as a *primary producer* is diminished, if a significant portfolio of mining innovation were developed leveraging the current intensity of primary production activity. Australia's mining software vendors already occupy an important niche in the global mining industry. There are

significant untapped opportunities for Australia to become a leading international provider of mining services, leveraging the service science innovations discussed above.

HEALTH AND AGED-CARE SERVICES

Australia's healthcare system is under severe pressure as it deals with a growing and aging population with resources that have not kept pace with demand. Efficient resource management is probably the biggest challenge for the healthcare and aged-care system. As with each of the three other sectors discussed above, service science research has the potential to deliver solutions for optimal system design (optimal collections of resources under appropriate configurations) as well optimized operations management. Consider, for instance, a specific component of the system, such as a hospital's radiotherapy department. Such a department would require radiation oncologists, radiation physicists, radiotherapy technicians as well as nurses and would operate under well-understood conditions of average demand (the number of new patient presenting over a given time period). For an administrator operating under severe resource constraints, a critical question to answer would be: what is the minimum staffing level that the department could be operated at, under assumptions of average load/demand, without having the system "break"? In this instance, the system would "break" if any patient incurred a wait for treatment – which patients under the circumstances can ill-afford. The solution relies on machinery for the optimal allocation of personnel and machines (specifically linear accelerators) to patients, and involves a combination of simulation, optimization and service monitoring technologies. In fact, the solution is remarkably similar to that for the generic *optimal service provisioning* problem discussed earlier.

Another version of the optimal service delivery problem is also relevant in this context. Consider a hospital chemotherapy department that has to treat a given (average) number of patients per day, with each patient having to follow a distinct chemotherapy protocol unique to the patient's condition. In effect, the protocol for each patient represents an instance of a particular clinical process design. An operational challenge is to identify an optimal allocation of nurses from a given pool of specialist chemotherapy nurses to each task in the pool of process instances (with one process instance for each patient at any point in time). Such problems pose important challenges for the service science community to extend and adapt its existing toolkit of techniques.

Quality assurance and monitoring of clinical processes represent another imperative. Healthcare decisions need to be logged and made available for auditing. The medical community has standardized *clinical guidelines*, which are in effect process descriptions for treatment processes for a range of conditions. In an idealized setting, every clinical process instance would be executed in conformance to the applicable clinical guideline. However, actual treatment instances routinely deviate from guidelines, simply because clinicians need to customize

guidelines to suit the particular circumstances of each patient. These raise important questions from the perspective of business process and service compliance.

Healthcare provides another important domain for the application of the socio-technical service systems frameworks discussed earlier. Healthcare processes have significant human-executed components (but also machine-executed components) which often display high degrees of variability, as discussed above. This poses unique challenges for the monitoring, quality assurance and general management of healthcare systems. Healthcare might well represent the upper-bound in terms of complexity of the socio-technical service systems that the community will have to understand and learn how to support.

BUILDING AN EXPORT-FOCUSED SERVICES SECTOR

Australia needs to develop an export-focused services sector (its higher education sector has had considerable success in this space, and the discussion below does not apply to it). Many fear that both manufacturing and services will progressively be off-shored, leaving Australia as a net consumer of both. Yet there are early, but tantalizing pointers to an emerging phenomenon that suggests that this need not be the case. Novel delegation relationships are being reported, where an Australian firm outsources some of its functions to an overseas firm which in turn sources some its functions from a (usually distinct) Australian firm. In other words, we are beginning to observe a geographic circularity in off-shoring relationships. The services being sourced from Australian firms were typically higher-end, knowledge-based services. This, perhaps, points a way forward for Australia to position itself as a global provider of higher-end services.

Australia, of course, will not be without competition as it seeks to achieve this position. Yet, Australia can differentiate itself from potential competitors in the following three ways. First, service science thinking could be leveraged to systematize and institutionalize innovation within Australian firms. Second, the technological and methodological insights into efficiency enhancement discussed above could be used to make Australian firms the most efficient service providers amongst their international peers. Third, Australian firms could leverage the insights into servitization discussed above to offer services in novel ways that make it easier for off-shore customers to monitor and quality assure both the processes and the deliverables.

THE NEXT STEPS

In terms of policy formulation, one might start with the standard prescriptions that documents such as this offer, such as greater engagement between the research providers and end-users and greater government funding. Both are important. The Service Science Society of Australia has engaged in a program of targeted sectoral workshops bringing together members of the services research community and key end-user representatives from the sectors discussed above. The Society will also continue to be pro-active in advocating the value proposition of this discipline and seeking greater engagement from government in this agenda.

It is imperative for the services community in Australia to go looking for problems of the kind outlined in this document. The results are likely to be extremely rewarding for the researchers involved.

The Society has laid the groundwork for its *Empirical Service Insights Initiative* which seeks to instrument an infrastructure for gathering service delivery data, correlated with service designs across all of the sectors discussed above. The data collected would provide a rich resource of generating insights on which service designs work well and why, which modes of service delivery are optimal, which methodological frameworks for service management work better in which contexts, and so on. If appropriately resources and executed, this could be an international trend-setter in terms of establishing a rich services research data infrastructure and repository.

CONCLUSION

In terms of the four questions posed at the beginning of this document, the following answers are offered. First, the services research community in Australia could do more to meet the needs of Australian society. Second, there is indeed a uniquely Australian services research agenda. While the individual elements of that agenda are not unique, the mix, and the emphases are. Third, this document lays down a roadmap for the kinds of services research that will likely have the most impact on national well-being. Finally, the roadmap offers a set of concrete prescriptions for the community to implement.

There are several other important sectors and application domains that this roadmap does not address in detail. It does not address the needs of the agricultural sector. It does not address the question of climate change services. Australia is at the forefront of international efforts at carbon mitigation, both in terms of technological advances as well as legislative/regulatory regimes. Ubiquitous decision support frameworks for carbon mitigation are an important challenge for the services community. Leveraging services thinking in the higher education sector is not addressed. The question of how Australian thought leadership in services can contribute to better outcomes in service delivery to the “bottom of the pyramid”, both in Australia and elsewhere, is also not addressed. These and others will be addressed in subsequent extensions to this roadmap.