

Towards the Visualisation of Collaborative Policy Decomposition

John Keeney, Owen Conlan, Declan O’Sullivan, David Lewis, Vincent Wade
*Knowledge & Data Engineering Group (KDEG),
Department of Computer Science, Trinity College Dublin, Dublin, Ireland.*
{John.Keeney, Owen.Conlan, Declan.OSullivan, Dave.Lewis, Vincent.Wade } @cs.tcd.ie

Abstract

Policy-based management is typically used to manage large and complex systems, e.g. the management of telecommunications networks. When collaborating multi-disciplinary managers try to create policies from their own viewpoint they are hindered by a lack of a common vocabulary and an inability to understand how the complex system currently operates and how it is constrained. For example, modern telecoms networks must be constrained by many different types of policies, including: operational policies, legal policies, billing and provision policies etc.. As the policy-set grows, and more cross-domain managers write policies, the system becomes unwieldy, static and stale. This is due to a lack of overall understanding of how the system is currently operating in a constrained manner. We propose to empower collaborating managers by building user-centric semantically informed visualisations of the managed system’s dynamic state and constraints. These visualisations and control adaptation mechanisms will be tailored to each manager’s context and domain thereby facilitating ongoing cooperative management and communication between collaborating managers.

1. Introduction

The increasing complexity of modern computing systems, e.g. large networks, makes their operational administration expensive and error prone, which presents significant obstacles to their continued use. The aim of policy-based management is to relieve cognitive load on human administrators by allowing the system operation to be managed by high-level constraints, specified as governance policies [1], thereby reducing management costs and errors. We aim to facilitate the collaborative policy-based management process by ensuring that each manager is presented with an understandable user-centric visualisation of appropriate information and necessary control mechanisms.

For example, in a collaborative network management scenario, how can a legal expert ensure compliance with privacy regulations when they do not

know how to set up secure VPN connections and do not appreciate the associated operational requirements, or how can a customer relations manager ensure gold-standard bandwidth allocation for a customer when they are unable to configure a network switch?

A significant drawback of policy-based management for complex systems remains the lack of an automated mechanism to resolve what is meant by high-level objectives used in governance directives [2,3,4,5,6]. The issue remains how to support the semi-automated mapping of high-level goals onto low-level enforceable actions, especially if the goals or actions are constrained or conflicting in some way and require a trade-off [4,5,7]. Any complex system will have domain experts at the “coalface”, who are very familiar with how the constituent parts of the system can be managed, and particularly aware of the operating constraints of those constituent parts. However, it is typically a team of people further up the management chain with differing abstract views of additional constraints who must actually decide how the composite managed system will be governed [1]. As the policy-set changes and grows, many of the management authorities are no longer able to understand how the managed system operates, much less how it is managed. Higher level managers may not be experts in other domains or about the intricate workings of the constituent components, and so they do not possess the knowledge or vocabulary to effectively collaborate with managers in different domains, so management cannot proceed effectively.

Human interactions with complex systems, particularly for management or tuning, tend to follow a pattern of *negotiation* [9], therefore the governance of such complex systems must be a dynamic, converging two-way process. Our research aims to facilitate this two-way process by visualising the constraints, state and context of the managed system in a dynamic user-centric manner, then allow the manager to manipulate the managed system at the same level of abstraction.

The benefit of presenting and comprehending large and complex data-sets in a visual manner is that it makes it much easier for the viewer to infer

conclusions from the knowledge presented [10,11]. However, in order to make sense of a complex task the user needs to be able to understand the information that informs the task and be able to abstract and contextualise this possibly unfamiliar information from a viewpoint that makes sense to them [11,12]. It is important that the user can see important and relevant information or controls to perform a management task, and that irrelevant information is hidden. However, the relevance of some information or mechanism is dependent on the context of the user and it must be contextualised for the user [12] in a way that the user appreciates its relevance. This has led to the popularity of personalised “dashboard” type monitoring applications for visualising key performance indicators (KPI) of managed systems [13], particularly in the area in network management. The personalisation of visual state and context representations is most important where visualisations are used as a communication tool between different knowledge domains [14].

In addition to visualisation, it is also necessary for the manager to manipulate and manage the system at the same level of abstraction and using the same vocabulary being presented to and being used by the manager. A user-focussed two-way process, which goes beyond just monitoring the system, will clearly assist in the system’s ongoing collaborative management.

2. Research Questions

We assert that the best way for cross-domain managers to collaboratively govern a complex system using policy-based management is for the state, constraints and policies of the system to be presented in a visual fashion, thereby relaxing the requirement for a common vocabulary [12,15]. In addition, this visualisation and associated control mechanisms must be personalised to each user’s context and requirements, so that what they see is tailored to their domain so they can understand it [12,14]. As collaborating managers change the system, each will see the composite effects of those changes from their own viewpoint. This raises a number of questions with respect to how to:

- present dynamic composite state, context and constraints from the many parts of the composite managed system to the manager, in a visual manner that is tailored to their priorities and context?
- allow collaborating managers to manipulate the managed system in a manner abstracted to their domain, and have these abstract management policies mapped to implementable actions?
- support dynamic feedback where managers can then see the effects of their actions in causally connected manner?

3. Ongoing Work

To support the answering of these questions this research is building on ongoing research in KDEG in TCD by exploiting ideas and mechanisms from the research domains of information visualisation, personalisation and semantically driven policy refinement approaches. As shown in figure 1 this framework is made up a number of sub-components. These include support for modelling the managed systems as a set of interoperating adaptive services that are manageable in a policy based manner, support for user-based semantic abstraction and presentation of state and constraints, and support for user-based policy decomposition to manipulate the managed system.

A. User-driven semantic abstraction of state and constraints

A Semantic Abstraction Layer (SAL) is under development which enables the markup, merging, and delivery of data from multiple heterogeneous data sources to applications which act as consumers of semantically annotated data [16,17]. In its first “data-centric phase” status information is received from domain-specific heterogeneous sources and is then annotated and semantically enhanced according to a domain specific ontology. The second “user centric phase” merges, cross-references and associates these annotated facts together depending on the preferences of the user. In this phase different priorities, translations and filters can be applied to the semantically rich information. This framework will need to be extended in order to incorporate a flexible model of constraints and policies to add to the data coming from the managed system.

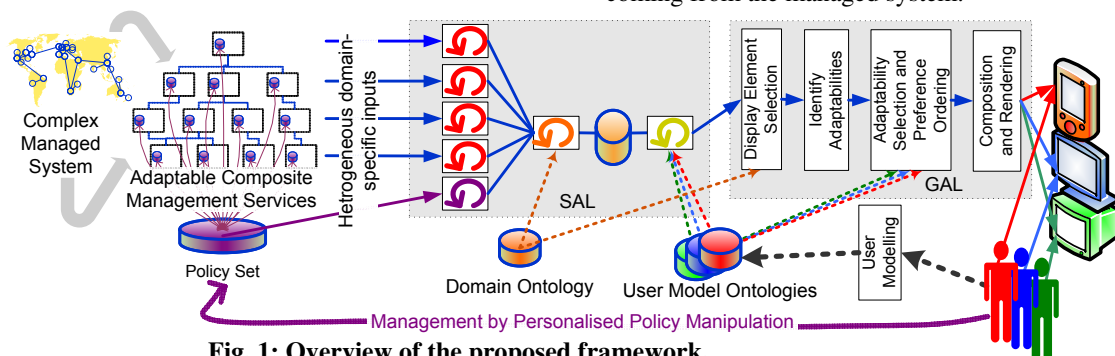


Fig. 1: Overview of the proposed framework.

B. Personalised visualisation of state and constraints

A Graphical Abstraction Layer (GAL) is also under development, which takes SAL output and processes it for user driven visualisations. First it maps the SAL outputs, and their associated semantics, to displayable elements and a set of scales of adaptability for each element. These adaptable displayable elements are then grouped, sorted and reworked according to user driven quantisation and interest level ordering, constrained by a model of the user. These tuned displayable elements are then composed for rendering according to the user's constraints. The GAL will require significant extensions with new adaptable display elements for this work. An extensive usability study will also be required to ensure that the users' expectations are being met.

C. Abstract manipulation of the managed system in a causally connected manner

Once the state and constraints of the system have been merged and translated to the context and domain of the manager then the manager will not be able to directly manipulate the system because the actual adaptability of the system will not be comprehensible without being itself translated for the manager. However, if the low-level enforceable concrete management actions are encoded in a semantically rich manner, then more abstract and higher level constraints can be composed from them in a way that they can be easily decomposed later [6,18]. These more abstract composed management actions can then be presented to the managers in a manner adapted to their context and knowledge domains. Without these semantics higher level constraints cannot be later translated and automatically decomposed into constituent low level enforceable constraints.

While there has been a body of research focussed on the presentation of domain independent policies, rules and constraints in a graphical manner (e.g. [4,20]), there has been little focus on presenting these visualisations as part of user-focussed control loop that incorporates monitoring and manipulation. The SAL and GAL frameworks mentioned are not currently focussed on forming part of a feedback control loop, and so do not address how the controlled system, i.e. the data source for the SAL, can be manipulated in a causally connected way. We need to extend our current SAL and GAL components by building upon ongoing research on semantic service techniques for policy-based object management [18,21]. These semantic service techniques have already been applied to network element management but can be applied to generic managed objects. Ensuring stability of this dynamic feedback control-loop is a key challenge since the user's feedback must allow the composite managed

system to converge to a steady state [23,24].

To address the we will continue work on modelling adaptive managed resources as adaptive services [18,21,25,26], which can be composed using existing service composition techniques to form a possibly very complex composite managed system. For this we plan to extend a policy refinement approach in [18,25]. High-level policies for a composite service, specified as finite state machine (FSM) transitions, can be automatically generated from state transitions for sub-FSMs describing the constituent services' manageable adaptive behaviours. This provides a two-way mechanism to map from composite policies to specific policies, and vice versa. However, such a discrete FSM-based model is not yet sufficient for complex composite services. A more expressive approach is planned where adaptivity is constrained so that the composed service operates within an acceptable behaviour envelope within a behaviour space [27].

Semantically rich high-level management policy goals can then be mapped and contextualised in a way similar to that which merges and translates the state and constraints of the system for visualisation. This contextualised service-oriented approach then provides a mechanism for the user to manipulate the managed object at a level of abstraction they can understand. Combined with a personalised dynamic view of the state and constraints of the complex managed system, the manager can then intuitively manage the system in a personalised service-oriented manner.

D. User modelling

In terms of populating a model of the user, much of KDEG's research on user modelling has been focussed on modelling learners to support e-learning environments [28, 31]. Many of these approaches can be exploited here to compile the user model required to drive the proposed semantic abstraction process to compose the adaptive visualisations and controls for the user. We plan to apply an ongoing user-modelling approach, whereby users specify their preferences and intent dynamically, thereby building up their user model in an ongoing manner. A framework developed by KDEG enables users to quickly set preferences in a highly visual and controlled manner [31]. The user model established using a mechanism similar to [31] will directly influence the personalisation available in the SAL, GAL and management feedback mechanisms.

4. Conclusions

Taken together the proposed framework (Fig. 1) will realise a two-way governance relationship, resulting in a truly user-centric approach to managing highly complex computing systems. It will present a way to

view the merged and abstract state and constraints of the composite managed system, at the level of abstraction that the collaborating manager can understand. The framework will surpass current monitoring tools (e.g. KPI “dashboards”) to provide a way to manipulate the composite system, at the same level of abstraction, and have these actions automatically mapped down to specific implementable changes. Additionally once management changes are made, the user will see the effects of these changes, in a causally connected way, in their own context. This user-focussed two-way process, which goes beyond just monitoring the system, will inform and facilitate feedback in the governance control loop and will clearly assist in the ongoing collaborative management of complex managed systems.

References

- [1] Barrett, R., “People and Policies: Transforming the Human-Computer Partnership”, Workshop on Policies for Distributed Systems and Networks (POLICY’04), New York, NY, 7-9 Jun 2004.
- [2] Sloman, M., “Policy Driven Management for Distributed Systems”, *Journal of Network and Systems Management*, 2(4), 1994.
- [3] “A Roadmap to Policy Research: Accomplishments and Future Challenges”, Expert panel discussion at Workshop on Policies for Distributed Systems and Networks (POLICY’05), Stockholm, Sweden, 6-8 Jun 2005.
- [4] Bandara, A., Lupu, E., Moffett, J., Russo, A., “A Goal-based Approach to Policy Refinement”, Workshop on Policies for Distributed Systems and Networks (POLICY’04), New York, NY, 7-9 June 2004.
- [5] Beigi, M., Calo, S., Verma, D., “Policy Transformation Techniques in Policy-based Systems Management”, Workshop on Policies for Distributed Systems and Networks (POLICY’04), New York, NY, 7-9 June 2004.
- [6] Rubio-Loyola, J., Serrat, J., Charalambides, M., Flegkas, P., Pavlou, G., “A methodological approach toward the refinement problem in policy-based management systems”, *IEEE Communications*, 44(10), Oct. 2006.
- [7] Lupu, E., Sloman, M., “Conflicts in Policy-Based Distributed Systems Management”, *IEEE Transactions on software engineering*, 25(6), 1999.
- [9] Maglio, P. Campbell S. Kandogan E. “On the Need for Negotiation in Policy-based Interaction”, Conference on Autonomic Computing (ICAC 2005), Seattle, WA, Jun 2005
- [10] González, V. and Kobsa, A. “Benefits of Information Visualization Systems for Administrative Data Analysts”. Conference on Information Visualization (IV 2003). London, UK, July 2003.
- [11] Simon, H. A., & Larkin, J. H. “Why a diagram is (sometimes) worth ten thousand words”, *Cognitive Science*, vol. 11, 1987.
- [12] Novak, J., "Helping Knowledge Cross Boundaries: Using Knowledge Visualization to Support Cross-Community Sensemaking", Conference on System Sciences (HICSS-40), Hawaii, January 2007.
- [13] Palpanas, T., Chowdhary, P., Mihaila, G.A., Pinel, F., “Integrated model-driven dashboard development” *Journal of Information Systems Frontiers*, 9(2-3). Jul 2007
- [14] Burkhard, R. A., “Learning from Architects: The Difference between Knowledge Visualization and Information Visualization”, Conference on Information Visualisation (IV 2004), London, UK, July 2004.
- [15] Klein, S. “Knowledge Visualization in Practice: Challenges for Future Corporate Communication”, Conference on Information Visualisation (IV 2005), London, UK, Jul 2005.
- [16] Hampson, C., “Semantically Holistic and Personalized Views Across Heterogeneous Information Sources”, Workshop on Semantic Media Adaptation and Personalization (SMAP07), London, UK. 17-18 Dec 2007.
- [17] Hampson, C., Lewis, D., O’Sullivan, D., “VABUCE: The Visualisation of Adaptive Behaviour in Ubiquitous Computing Experiments”, In Proc. of Information Technology and Telecommunications Conference, Carlow, Ireland. 25-26 Oct 2006.
- [18] Keeney, J., Carey, K., Lewis, D., O’Sullivan, D., Wade, V., "Ontology-based Semantics for Composable Autonomic Elements", Workshop on AI in Autonomic Communications at IJCAI 2005, Edinburgh, Scotland. 30 Jul 2005.
- [20] Feeney, K., Tsarouchis, C., Lewis, D., “Policies as Signals in Collaborative Policy Engineering”, Workshop on Policy-Based Autonomic Computing (PBAC 2007), at ICAC 2007, Jacksonville, Florida, 15 Jun 2007
- [21] Lewis, D., Keeney, J., O’Sullivan, D., "Policy-based Management for Resource-Specific Semantic Services", Workshop on Distributed Autonomous Network Management Systems (DANMS 2006). Dublin. 16 Jun 2006.
- [23] Diao, Y., Hellerstein, J.L.; Parekh, S.; Griffith, R.; Kaiser, G.E.; Phung, D., "A control theory foundation for self-managing computing systems," *IEEE Journal on Selected Areas in Communications*, 23(12), Dec 2005.
- [24] Hellerstein, J.L., "Challenges in control engineering of computing systems," American Control Conference 2004, Boston, MA. 30 Jun-2 Jul 2004
- [25] Carey, K., Lewis, D., Higel, S., Wade, V., “Adaptive Composite Service Plans for Ubiquitous Computing”. Workshop on Managing Ubiquitous Communications and Services (MUCS2004), Dublin, Ireland. Nov 2004.
- [26] Lewis, D., Feeney, K., Carey, K., Tiropanis, T., Courtenage, S., “Semantic-based Policy Engineering for Autonomic Systems”, Workshop on Autonomic Communication, Berlin, Germany. 18-19 Oct 2004.
- [27] Dobson, S. and Nixon, P., “More principled design of pervasive computing systems”. Engineering for Human-Computer Interaction and Design, Specification and Verification of Interactive Systems. Hamburg. Jul 2004.
- [28] Conlan, O., Hampson, C., O’Keeffe, I., Heller, J., “Using Knowledge Space Theory to support Learner Modeling and Personalization”, World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education, Hawaii, USA, 2006.
- [31] Spath, M., “A graphical User interest and preference modelling tool exemplified on Sightseeing planner for Dublin”, MSc Thesis, Computer Science, Trinity College Dublin, Ireland. Sept 2007.