Potential Benefits and Challenges of Closed-Loop Optimization Processes for IT Support Organizations

Claudio Bartolini
HP Labs
Hewlett Packard
Palo Alto, CA, USA
claudio.bartolini@hp.com

Cesare Stefanelli, Mauro Tortonesi
Department of Engineering
University of Ferrara
Ferrara, Italy
{cesare.stefanelli,mauro.tortonesi}@unife.it

***Abstract*—IT services are getting increasingly complicated, and require IT support organization to manage them. IT support organization are in charge of the incident management process and represent mission-critical structures whose performance needs to be frequently assessed and optimized. State-of-the-art research in the performance optimization of IT support organization proposes user-driven performance assessment and optimization processes based on what-if scenario analysis tools that implement sophisticated IT support organization models. This manuscript instead represents a preliminary study of a different kind of optimization processes, of the closed-loop type, that try to autonomously identify optimal IT support organization configurations according to inputs provided by the user. This paper discusses the development challenges in realizing decision support tools for closed-loop optimization processes and presents a prototype system. The preliminary evaluation of our tool demonstrates that closed-loop processes might be impractical as reference tools but can effectively complement and extend human-driven ones.**

Decision support, Information Technology Infrastructure Library (ITIL), IT service management, incident management.

I. Introduction

IT services are getting more and more complicated, and their management represents a major challenge. IT support organizations are the entities in charge of the incident management process, as defined by ITIL [1] [2], and their optimization is of critical importance.

While implementing a relatively simple business process, IT support organizations are complex systems composed of support groups that cooperate, often through tightly coupled interactions, to restore normal service operations after a disruption. The performance optimization of IT support organization is therefore a considerably challenging task.

In our previous works, we proposed and advocated an interactive and user-guided optimization process for IT support organizations [3] [4]. Expert human decision makers are in charge of guiding the optimization process towards the desired performance objectives, which are organization-specific and defined by the business management.

The present work, instead, takes an exploratory approach in the opposite direction and attempts to study the potential benefits that an automated, closed-loop optimization process would bring, as well as the fundamental challenges that its development would pose. Closed-loop optimization processes try to identify the optimal IT support organization configuration according to inputs provided by the user: an IT support organization model, a series of constraints on possible reconfigurations, and a performance evaluation criterion.

This paper discusses the challenges in the development of closed-loop optimization processes. To evaluate the potential benefits that closed-loop optimization processes could bring, we have realized a prototype decision support tool that leverages on what-if scenario analysis and extensive user-provided input to identify an IT support organization configuration with an optimal business impact.

While it builds on extensive research work that the authors have dedicated to IT support organization modeling, this paper presents 2 important original contributions: a method to rigorously define the space of possible IT support organization configurations and an extended business-impact model that enables the performance assessment of the IT support organization according.

Preliminary evaluations of our tool demonstrate that, while the adoption of closed-loop performance optimization as a reference decision support tool might be impractical, these processes can complement and extend interactive, human-driven optimization processes.

II. IT Support Organizations

IT support organizations are composed of a possibly large number of operators, with different specialization, skills and work shift schedules. Operators are organized in support groups of different levels (usually three to five), with lower level groups managing generic issues and higher level ones handling mundane and time-consuming tasks. The Help Desk is the lowest support group that interfaces with customers reporting an IT service disruption, and consequently "opens" an incident, sometimes also called trouble-ticket or simply ticket. The incident is then "assigned" to a specific support group for its solution, and can possibly escalates to a higher level group in case of a technically deeper support needed.

During its lifetime, an incident moves through different states and is possibly managed by several groups of different levels. In each state, the incident record is updated with pertinent information, such as state name and related service restoration activity. If the customer request the organization to stop working on the incident, the incident is placed in a "suspended" state to avoid incurring into Service Level Objective penalties. The ticket is placed in the "closed" state when the disruption is repaired, and it is "resolved" when the customer confirms that the service has been fully restored.

In some ways, IT support organizations resembles telephone call centers [5] [6] and can thus be modeled as open queuing networks [7]. However, notice that IT support organizations have some peculiar characteristics that distinguish them from telephone call centers. In fact IT support organizations do not need to consider call blockings, abandonments, or redials, but should instead consider the routing of incidents through the system as well as complex incident prioritization policies, since IT support organizations might serve many customers with different profiles, each one with a specific SLA.

III. AUTOMATED PERFORMANCE OPTIMIZATION OF IT SUPPORT ORGANIZATIONS

When dealing with the performance optimization of IT support organization, the state-of-the-art of research is represented by decision support tools that enable a user-driven process based on purposely developed models of IT support organizations and what-if scenario analysis [3]. These tools reenact the behavior of real-life IT support organizations through simulation techniques and enable an iterative performance optimization process. As depicted in Fig. 1, users can optimize the performance of their IT support organizations by incrementally specifying the set of changes to apply to the current organization configuration in order to define an alternative configuration that will be tested on a set of performance metrics.

The inclusion of a human expert in the optimization loop presents significant benefits, mostly because putting a human in charge of decision making enables him to develop a deeper knowledge of the system as well as of the optimization objectives through an interactive learning process. In fact, a human decision maker might improve the quality of his decisions by analyzing the impact that past decisions had on the IT support organization performance, and possibly changing his decision making criteria and/or his target optimization objectives along the optimization process. (This essentially warrants that the optimization process stops only when it leads to acceptable results.)

In addition, humans can also consider other factors in their decision making that would be difficult to model by what-if scenario tools, such as outsourcing or other reorganization opportunities.

Finally, as decision making is taken care of by a human, there is no need to develop sophisticated multiple-criteria decision making software systems to comparatively evaluate the performance of different IT support organization configurations [8].

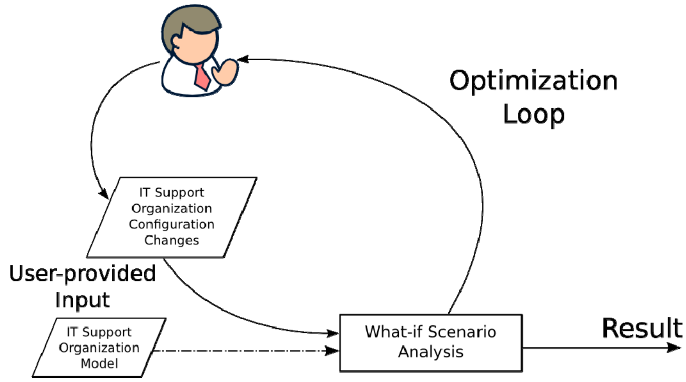

Figure 1. Workflow for the user-driven performance optimization processes.

However, human-driven performance optimization processes specifically require the inclusion of an expert IT manager (possibly with an in-depth knowledge of the IT support organization modeled by the tool) in the optimization loop and can be very time-consuming and. In addition, human-driven processes, that typically consist of a series of small performance improvement steps, might converge to local optima, thus precluding further potential performance optimization opportunities. Finally, human decision makers are often irrational and error prone, thus requiring specifically developed decision support systems in order to mitigate the likelihood of achieving sub-optimal results in performance optimization [9] [10].

These considerations suggest to consider automated, closed-loop performance optimization processes to support and complement human-driven ones. By leveraging on the same IT support organization models and what-if scenario analysis techniques discussed above, automated performance optimization processes have the benefit of autonomously exploring the space of possible configurations. Automated tools can be instructed to explore many avenues of optimization, thus minimizing the risk of terminating the optimization process when a local optimum is reached.

However, the realization of automated processes for the performance optimization of IT support organization presents several challenges. First, it requires a mechanism that enables the user to define a set of constraints that delimit the space of alternative configurations to consider, as well as to develop procedures that explores the space within the above mentioned constraints. In addition, the process must be capable of reenacting IT support organizations' behavior with an alternative configurations. Finally, the process needs a procedure that measures the performance of the IT support organization with the new configuration, enabling to compare it with that of other configurations.

Evaluating the performance of an IT support organization configuration is a very difficult task. In fact, the full extent of alternative configurations of IT support organizations cannot be captured through IT-level metrics alone. In addition, IT-level metrics are plenty and the use of automated performance optimization process would force to consider multi-criteria decision making methods.

Instead, *business impact analysis* represents a significantly better criterion to adopt for the performance optimization of IT support organizations. Business impact-driven optimization aims at minimizing the adverse impact of service disruptions on the business, by considering all the costs attached to critical incident occurrences. In addition, business impact represents a single measure, that therefore significantly simplifies the comparison process, as there is no need to consider multi-objective optimization methods.

From a theoretical perspective, the performance optimization procedure thus becomes a nonlinear programming problem, of the following kind:

$$\begin{aligned} &\min BI(x) \\ &\textit{subject to } x \in S_{PC} \end{aligned} \quad (1)$$

where the variable x represents the IT support organization configuration; the set S_{PC} represents the space of possible IT support organization configurations to explore; and the function BI represents the performance of the IT support organization with configuration x, i.e., its business impact. The optimization process starts from x_0, the IT support organization configuration currently in place.

To study the challenges involved in the development of closed-loop performance optimization process in practice, as well as their potential benefits, we have realized a prototype decision support tool. Our system leverages on what-if scenario analysis and extensive user-provided input to identify an IT support organization configuration with an optimal business impact. The workflow for the automated optimization process implemented in our prototype tool is depicted in Fig. 2.

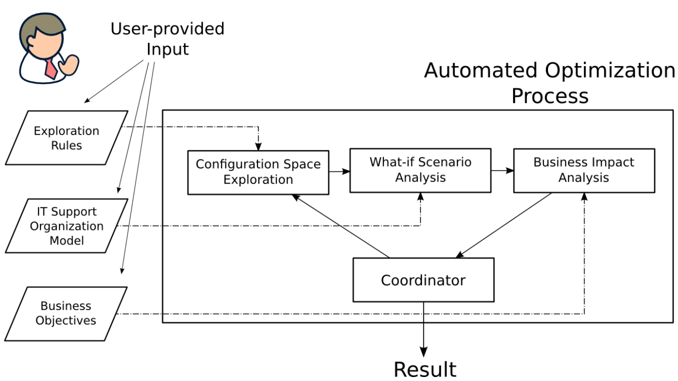

Figure 2. Workflow for the automated performance optimization processes.

IV. MODELING IT SUPPORT ORGANIZATIONS

With regards to the modeling of IT support organizations and to the evaluation of alternative configurations for the organization modeled, the procedure presented in this paper leverages on the extensive experience and tools that the authors developed in the Symian research project.

Symian is a decision support tool that allows to create realistic models of real-life IT support organizations and to reenact them, as well as modified versions, to fully understand the impact of IT processes, strategies and tactics on the organizations' performance.

Symian enables to play what-if scenario analysis by exploiting a discrete event simulator to reproduce in detail the behavior of IT support organizations and to evaluate their performance in managing incidents. The simulation approach is particularly appropriate, given that the scale and the complexity of real life organizations make it extremely difficult to devise an analytical model.

Symian proposes a model of IT support organizations based on open queuing networks [3]. More specifically, Symian models IT support organizations as a set of interconnected support groups that exchange tickets according to a Markovian routing process defined by a stochastic transition matrix. In turn, support groups are modeled as multi-server *first-come-first-served (FCFS)* queues. (The authors recently developed a significantly more sophisticated model of support groups based on *multiple-priority* queues, that will not be considered in this paper for the sake of simplicity. The interested reader is referred to [11].)

Open queuing networks are particularly attractive for the reenactment of IT support organizations, as they represent a very good tradeoff between model complexity and model accuracy. In fact, they can easily measure IT support organization dynamics in terms of throughput, queue lengths, response times, and utilization, both at the system level and at the single support group level. Our model proved capable of reenacting the behavior of real-life enterprise-class IT support organizations very well, both at the system-wide and at the single support group level.

To accurately capture a wide range of possible behaviors, our model reenacts the behavior of a support group according to a large set of model parameters, such as a *service time distribution*, a *set of operators*, and configurations such as a *priority assignment policy*, an *operator assignment policy*, and the operators' attributes *(workflow, skill set, etc.)*.

Fig. 3 provides pictorial representations of how our FCFS and multiple-priority model works. When an incident arrives to a support group, it is attached a service time, randomly sampled from the service time distribution associated with the corresponding support group. Service time is an attribute that represents the amount of time that operators have to spend on the incident to finish the portion of work that competes to the current support group. Once assigned a service time, the incident is then inserted at the end of the corresponding queue.

As one or more operators become available to start servicing new incidents from the queue, our model starts assigning them the tickets from the beginning of the queue. The incident-operator association procedure is performed according to the configured operator assignment policy, which specifies the set of rules to use for selecting which incident an operator should start working on when he becomes idle. The operator assignment policy is probably the most important configuration parameter of our model, as it can have a very large impact on the support group behavior. The simplest supported operator assignment policy assigns incidents to idle operators of the corresponding group in a round-robin fashion. However, it is also possible to reenact more sophisticated behaviors. For instance, for a multiple-priority support group model one could configure an operator assignment policy that at the arrival of incidents of a specific type, .e.g., with higher priority, preempts operators servicing lower priority ones, or that enforces the servicing of some particularly complicated incidents only by highly skilled and/or experienced operators.

Our model allows for very sophisticated operator management. In fact, the model provides support for operator work shift configuration, enabling an extremely realistic representation of the support group behavior. More specifically, the implementation of our model leverages on a simulated clock that reenacts the flow of simulation-time in a very similar way to what happens in real life. If an operator's work shift ends before the incident service time is expired, incidents can either be handed over to another operator for around-the-clock servicing or simply wait until the operator's next work shift. In addition, as in real life IT support organizations it is not uncommon for operators to work on multiple tickets at the same time, our model also supports operators that simultaneously work on more than one incident. Finally, our model can be configured to consider specific operator skills that skew their ability to deal with incidents of a particular category. An operator with a 2.0 skill parameter in a specific incident category will take twice as less time to resolve an incident of that category than a normal operator.

When its associated service time is exhausted, the incident can leave the support group. In case further work is required to resolve the related service disruption, the incident will be escalated to a different support group. Else, the incident is considered resolved and its lifecycle ends.

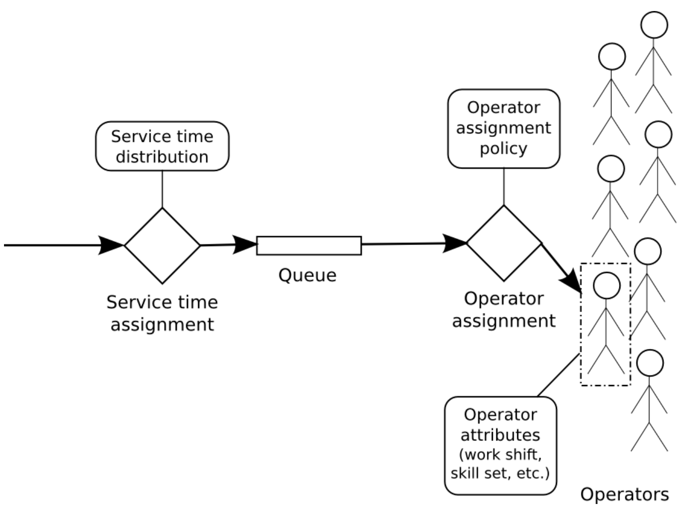

Figure 3. FCFS queuing model for IT support groups.

V. DEFINING AND EXPLORING A SPACE OF POSSIBLE CONFIGURATIONS FOR AN IT SUPPORT ORGANIZATION

The first step of a closed-loop automated IT support organization performance optimization process would be the definition of a space S_{PC} of alternative configurations to consider as well as a strategy to explore the space within the above mentioned constraints.

Notice that S_{PC} is a subspace of S, the multidimensional space of all physically possible IT support organization configurations. Therefore, the definition of S_{PC} requires to apply constraint to S. However, users cannot be expected to express these constraints as hyperplanes or convex functions in S. Instead, we need to consider rules such as "can decrease service time up to 5% in all support groups" and "can transfer up to 50% of the operators from support X to support group Y". This complicates optimization problem (1) and requires explicit support for domain-specific solutions that enable the definition of constraints in S.

To support these steps, we have designed a system that enables users to express constraints for the exploration of the IT support organization configuration space to consider by specifying easily understandable rules. Our system proposes 2 basic types of rules: the *explore parameter range* rule and the *add/remove/transfer asset* rule. We have also developed a domain-specific language that enables users to define constraints through the above mentioned rules. Fig. 4 shows an example of the constraints that our tool enables to specify.

Through rules of the explore parameter range type, the user defines its interest to explore a given area of the parameter space. Rules of the asset type instead enable the use to state his interests in exploring what would happen by adding some assets, e.g., operators, to a support group, or removing them, or transferring them to another group. For instance, rule #1 in Fig. 4 states the user's intention to explore alternative configuration with a different mean service time parameter up to a 5% relative change.

Instead, through rules of the move asset type, a user states his intention of exploring different configurations that transfer an asset from a support group to another. Rule #2 in Fig. 4 states the user's intention to explore alternative configuration where a number of operators from support group SG1 are transferred, possibly in retraining, to support group SG2.

Finally, let us point out that from a mathematical point of view, an explore parameter range rule translates into a hyperbox-type constraint within the space S of all possible IT support organization alternative configurations. Asset type rules, instead, translate into a hyperline-type constraint in S. The combination of these rules allows to define highly non-linear spaces S_{PC}'s, thus forcing to consider very robust space exploration methods.

After the definition of the space constraints, a proper space exploration strategy needs to be selected. Since the function that the tool needs to optimize is highly non-linear and there are no information about its derivatives, we cannot leverage on gradient-descent based optimization techniques such as the Broyden–Fletcher–Goldfarb–Shanno (BFGS) [12] and related methods. Instead, metaheuristics represent better suited solutions. In addition, since the space to explore is typically wide, we need to consider metaheuristics designed for large-scale optimization.

In addition, for rules of the add/remove asset or transfer types there is the problem of limiting the number of configuration changes to consider. This forces to use space exploration strategies that consider paths of limited length when exploring the configuration space, possibly keeping a memory of past changes to improve the exploration.

To address these issues, there is the need to develop tools that supports several exploration strategies. So far, we have mostly experimented with metaheuristics based on stochastic algorithms, such as random search, adaptive random search, and Tabu search, and implemented them in our tool. Genetic algorithms also represent very promising exploration strategies.

```
explore_range :mean_service_time,
       :upto => 5.percent
move_asset :operators, :up_to => 10,
       :from => "SG01", :to => "SG22"
exploration_strategy :random
```

Figure 4. Example of constraint definition on the multidimensional space of possible IT support organization configuration parameters.

VI. BUSINESS-IMPACT EVALUATION OF ALTERNATIVE IT SUPPORT ORGANIZATION CONFIGURATIONS

This section presents a business impact model, an extended version of the one proposed in [13]. In the Business Impact Analysis phase, our tool evaluates the performance of the IT support organization configuration according to business impact criteria.

To this end, the tool leverages on user-provided business management preferences. More specifically our automated optimization process computes the business impact on the given IT support configuration x by estimating its *alignment* with the user-provided business objectives.

Business objectives, defined at the business management level, are the targets of the incident management process. Our tool models business objectives as a *target region* (usually a range of values) over a *KPI* (Key Performance Indicator). If a KPI value lies within the corresponding target region at the end of the evaluation period, then the business objective is met.

Leveraging from the Balanced Scorecard methodology [14], we group business objectives in perspectives and associates to each perspective a *weight*, which captures the importance of the objective and is defined by the user. Within every perspective, each single business objective is also assigned an importance weight.

TABLE I. EXAMPLE OF BUSINESS OBJECTIVES.

Business objective	**KPI**	**Target region**	**Obj. Wt.**	**Final Wt.**
Financial Perspective – Importance weight: 0.6				
Reorganization costs	Total cost of implementing new organization	lower than 15,000 $ per month	0.7	0.42
Costs for SLO penalties	Total Cost of SLO penalties	lower than 10,000 $ per month	0.3	0.18
External Perspective – Importance weight: 0.4				
Customer Satisfaction	Number of SLO violations	5 or less	1.0	0.40

To evaluate financial objectives, our tool leverages on a dedicated component that calculates financial costs. This procedure is divided in 3 steps or phases: *reorganization cost estimation*, *Service Level Objective (SLO) penalty estimation*, and *operational cost changes estimation*.

The reorganization cost estimation phase calculates the cost of implementing the given IT support organization configuration x with respect to the configuration x_0 currently in place. The reorganization costs directly related to the strategies implementation depend from the specific configuration x to consider. For instance, the cost of adding new operators to a specific support group must consider the costs for operator training, equipment, and salary. The cost of implementing software/hardware replacement and/or upgrades, instead, must consider the cost for buying new software/hardware and the cost for installation, configuration, and training. As a result, our tools requires the user to provide specific implementation costs for each strategy to be evaluated.

The second phase estimates SLO violation penalties due to strategies implementation. To this end, our tool requires users to define the conditions in which SLO violations occur and their penalty amount. The tool then uses the values of service level indicators obtained from the previous phase to find whether SLO violations occur in the context of the strategies under evaluation, and calculates their costs.

In the last phase, the tool considers the drift in operation efficiency as a cost to the business. In fact, if a different IT support organization configuration changes the throughput of ticket resolutions, that has an impact on the business. To take this into account, the model leverages on a monetization function that links the MICD (Mean Incidents Closed Daily) metric with additional efficiency-related costs. The monetization function, that must be provided by the user, will typically have an *arctg* or sigmoidal shape. This operation is depicted in Fig. 5.

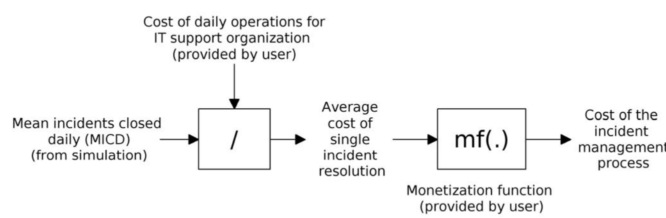

Figure 5. Calculating additional costs of the incident management process related to changes in the IT support organization performance.

VII. EXPERIMENTAL EVALUATION

To evaluate the closed-loop process, we applied it on a model of real-life IT support organization that we build by analyzing transactional log data provided to us by the Outsourcing Services Division of Hewlett-Packard. HP Outsourcing manages, among other IT services, the Help Desk function on behalf of various enterprise customers. The data used for this experiment comes from the subset of the organization serving a single enterprise customer from the financial services industry, whose name will be disguised as BailUsOut in the remainder of the paper.

The BailUsOut IT support organization is run by the Outsourcing Services Division of HP. Having a global 24/7 presence, HP Outsourcing faces the daily challenge of supporting multiple environments for multiple customers in disparate geographies. Hundreds of support groups employing thousands of engineers provide support to clients all around the world. In particular, 34 support groups are dedicated to BailUsOut, and 38 more groups have shared responsibilities across multiple enterprise customers and deal with tickets generated by BailUsOut. These support groups are geographically distributed and each work in their own local time-zone.

We present an illustrative example that demonstrates the potential that the autonomous exploration of the configuration space performed by closed-loop processes has. To this end, we adopted a very simple business impact model, that considers only costs of implementing new IT support organization, and a simple exploration strategy depicted in Fig. 6.

```
explore_range :mean_service_time,
      :upto => 100.percent
exploration_strategy :random, :runs => 30
```

Figure 6. Exploration rules for the BailUsOut IT support organization model.

The results presented in Fig. 7, which compares business impact and IT level metrics obtained from the application of the closed-loop optimization process to the BailUsOut IT support organization, demonstrates that the autonomous exploration of the configuration space might hint to interesting optimization avenues that could be further explored in user-driven processes.

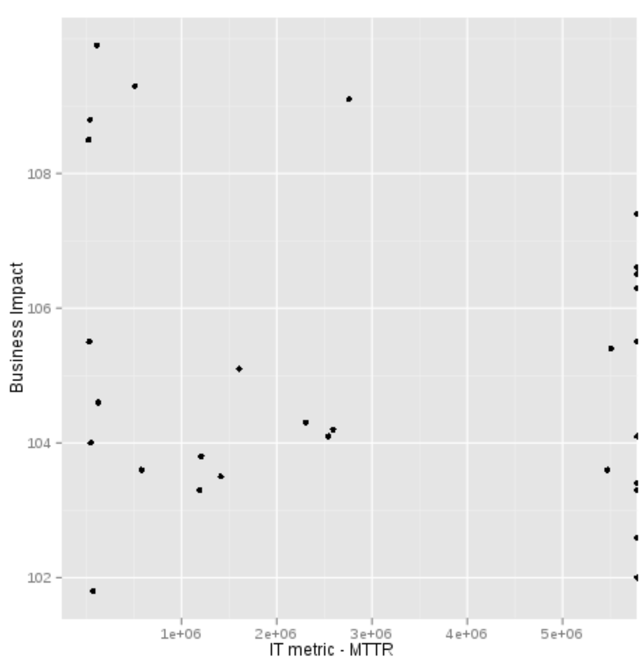

Figure 7. Comparison of business impact and IT level metrics from the closed-loop optimization of the BailUsOut IT support organization model.

VIII. Related Work

The present work belongs to the business operation analysis research area. More specifically, it falls within research approaches that aim at improving business processes through the collection of metrics on business processes and the inferences over the acquired data [15] or the exploitation of simulation methods [16]. However, most research studies in business operation analysis propose techniques tailored to complex process descriptions, consisting of many quite well defined steps with a limited number of alternative paths. Instead, incident management is a relatively simple process taking place in a complex organization, and can by characterized by a few simple steps with a huge fan-out of alternative possibilities. These peculiarities make common business process management techniques either overkill or unapplicable to the incident management process.

Other research efforts have focused on optimization of IT-specific business process. Among these, a particularly interesting series of studies is Diao et al.'s, which focuses on the estimation of labor cost and business value of IT services from the analysis of process complexity [17] [18]. The present work represents a significantly different approach from Diao et al.'s, because Symian-Web aims at improving IT support organizations' performance through decision support and simulation techniques.

Among other research efforts modeling IT support organizations , a particularly interesting one is Shao et al.'s [19] [20]. Their EasyTicket system aims at optimizing the ticket routing in IT support organizations through machine learning techniques, and is based on a queuing network-based model similar to the one adopted by Symian-Web's. EasyTicket is based on the assumption that each type of incident has a specific resolver support group, and that therefore by facilitating routing incidents to their resolver support group as quickly as possible the IT support organization performance could be improved. In our experience, this assumption does not always hold in practice, as usually in real-life IT support organizations several support group need to cooperate in order to restore major service disruptions. The Symian-Web approach, instead, does not rely on any restrictive assumption about the incident resolution process, thus enabling a comprehensive analysis of the organization performance which considers both incident routing effectiveness and efficiency within each individual support group.

IX. Conclusions and Future Work

The present work represents a first study of closed-loop optimization processes for IT support organization. Compared with the interactive/user-driven ones, closed-loop performance optimization processes require a substantially larger initial input from the user. This might make closed-loop processes impractical.

However, they presents the benefit of autonomously exploring the space of possibility. Hence, automated processes are probably more suited to complement user-driven interactive optimization processes, instead of replacing them.

REFERENCES

[1] Office of Government Commerce, "Service Strategy Book", The Stationery Office, Norwich , 2007.

[2] Office of Government Commerce, "Service Operation Book", The Stationery Office, Norwich, 2007.

[3] C. Bartolini, C.Stefanelli, M. Tortonesi, "SYMIAN: Analysis and Performance Improvement of the IT Incident Management Process", *IEEE Transactions on Network and Service Management*, Vol. 7, No. 3, pp. 132-144, September 2010.

[4] C. Bartolini, C. Stefanelli, M. Tortonesi, "Modeling IT Support Organizations Using Multiple-Priority Queues", accepted for publication in *Proceedings of the 13th IEEE/IFIP Network Operations and Management Symposium (NOMS 2012) - Main track*, 16-20 April 2012, Maui, Hawaii, USA.

[5] G. Koole, A. Mandelbaum, "Queueing Models of Call Centers: An Introduction", *Ann. of Oper. Res.*, Vol. 113, No. 1-4, pp. 41-59, July 2002.

[6] G. Koole, "Call Center Mathematics", VU University of Amsterdam, Department of Mathematics Website. http://www.gerkoole.com/CCO/. Accessed 9th January 2012.

[7] G. Bolch, S. Greiner, H. de Meer, K. Trivedi, "Queuing Networks and Markov Chains: Modeling and Performance Evaluation with Computer Science Applications'', 2nd Edition, Wiley 2006.

[8] J. Branke, K. Deb, K. Miettinen, R. Slowinski, "Multiobjective optimization: interactive and evolutionary approaches", Springer, 2008.

[9] D. Kahneman, P. Slovic, A. Tversky (Eds.), "Judgment Under Uncertainty: Heuristics and Biases", Cambridge University Press, 1982.

[10] T. Gilovich, D. Griffin, D. Kahneman (Eds.), "Heuristics and Biases: The Psychology of Intuitive Judgment", Cambridge University Press, 2002.

[11] C. Bartolini, C. Stefanelli, M. Tortonesi, "Modeling IT Support Organizations Using Multiple-Priority Queues", accepted for publication in *Proceedings of the 13th IEEE/IFIP Network Operations and Management Symposium (NOMS 2012) - Main track*, 16-20 April 2012, Maui, Hawaii, USA.

[12] J. Nocedal, S. Wright, "Numerical Optimization", 2nd Ed., Springer, 2008.

[13] C. Bartolini, C. Stefanelli, M. Tortonesi, "Business-impact analysis and simulation of critical incidents in IT service management", in *Proceedings of the 11th IFIP/IEEE International Symposium on Integrated Network Management (IM 2009)*, pp. 9-16, 1-5 June 2009, New York, NY, USA.

[14] R. S. Kaplan, D. P. Norton, "The Balanced Scorecard: Translating Strategy into Action", Harvard Business School Press, 1996.

[15] F. Casati, M. Castellanos, U. Dayal, M.C. Shan, "A Metric Definition, Computation, and Reporting Model for Business Operation Analysis", in *Proc. of Advances in Database Technology - EDBT 2006, 10th International Conference on Extending Database Technology*, Munich, Germany, March 26-31, 2006.

[16] K. Tumay, "Business process simulation", in *Proc. Winter Simulation Conference 1995*, 3-6 December 1995.

[17] Y. Diao, A. Keller, S. Parekh, V. Marinov, "Predicting Labor Cost through IT Management Complexity Metrics" in *Proceedings of the 10th IEEE/IFIP Symposium on Integrated Management (IM 2007),* Munich. Germany.

[18] Y. Diao, K. Bhattacharya, "Estimating Business Value of IT Services through Process Complexity Analysis", *in Proceedings of the 11th IEEE/IFIP Network Operation and Management Symposium (NOMS 2008)*, Salvador de Bahia, Brazil.

[19] Q. Shao, Y. Chen, S. Tao, X. Yan, N. Anerousis, "Efficient Ticket Routing by Resolution Sequence Mining", in *Proceedings of 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'08)*, Las Vegas, NV, USA, 24-27 August 2008.

[20] Q. Shao, Y. Chen, S. Tao, X. Yan, N. Anerousis, "EasyTicket: A Ticket Routing Recommendation Engine for Enterprise Problem Resolution", in *Proceedings of 34th International Conference on Very Large Data Bases (VLDB'08)*, Auckland, New Zealand, 23-28 August 2008.