White paper

Optimization of Modern Data Centers using Fujitsu's Quantum-Inspired Solution

Given the world's increasing dependency on IT, the complexity of data centers grows significantly. In this white paper, Fujitsu and the Magdeburg Research and Competence Cluster compare the new Fujitsu quantum-inspired solution with state-of-the-art solution algorithms in order to evaluate the potential of the Digital Annealer for server consolidation and optimization problems as they frequently appear as part of operational SAP IT infrastructure analysis.

Content

Enterprise Application Capacity Management 2
The Optimization Problem 3
The Digital Annealer: A Quantum-Inspired Computing Technology 4
Introduction of Solution Algorithms 5
Results of a Multi-Case Data Center Capacity Management Study 8
Conclusion 10
Executive Summary 10
Acknowledgement 11
Authors 11
References 11

Enterprise Application Capacity Management

Given the world's increasing dependency on IT, the complexity of data centers grows significantly. At the same time, costs to operate and maintain infrastructures are supposed to be reduced for example by optimizing the infrastructure landscape. According to best practices, capacity management is an essential IT process with the objective to ensure sufficient resource capacity for any given IT service as part of its design stage.

The Complexity of Server Consolidation

In a top down manner, IT services are to be aligned with ever-changing business requirements and, at the same time, must utilize hardware components efficiently. This challenge is addressed by server consolidation exercises which aim at balancing operational costs and non-functional requirements such as performance. Idling components represent saving potential that can be identified and addressed by optimization techniques in order to find an optimal or near-optimal distribution of IT services across the component layer. Therefore, server consolidation is a field with various challenges and thus a good opportunity to assess the capabilities of the Digital Annealer. Solution deployment is typically enabled by server or application virtualization techniques, as demonstrated in integrated systems such as Fujitsu PRIMEFLEX. During the design stage, the main challenge is to find feasible designs, which ensure not to overload servers at any time with respect to the dynamic resource demands of all running services. Since the number of services and servers generally increase over time, the solution space has dramatically grown to a size, which cannot be managed manually anymore. In fact, heuristics and metaheuristics are applied to complex consolidation problems in order to identify feasible solutions of this multi-dimensional problem. With the Digital Annealer, a new technique emerged to search through large solution spaces for optimal designs. This white paper compares this new quantum-inspired solution, developed by Fujitsu, with state-of-the-art solution algorithms, in order to evaluate the potential of the Digital Annealer for server consolidation and optimization problems as they frequently appear as part of Fujitsu's SAP SystemInspection Service.

Fujitsu SAP SystemInspection Service

Fujitsu offers a number of Technical Advisory Services for infrastructure solutions to help customers cope with today´s capacity management challenges for their SAP enterprise applications. The SAP SystemInspection Service is a unique offering in the field of infrastructure optimization as it relies on decades of practical experience combined with state-of-the-art research insights from an ongoing academic partnership. With a strong focus on SAP Technology such as SAP NetWeaver and SAP HANA, the SAP SystemInspection Service offers an efficient analysis and comprehensive consultation package for existing infrastructure environments at a fixed price. The goal is to get a complete understanding of the current workload, performance and related resources consumption and distribution in order to provide measures and guidance to optimize SAP Landscapes according to the business strategy and requirements. The use cases described in this white paper were conducted as part of Fujitsu's ongoing efforts to make scientific contributions and, finally, push forward Fujitsu's customer experience related to both SAP infrastructure management and optimization using cutting-edge quantum-inspired technology. More information about the SAP SystemInspection Service is available at the end of this white paper.

The Optimization Problem

The IT capacity management process aims at continuously balancing costs and performance. Since enterprise servers are often under-utilized, server consolidation is an effective means to address this trade-off. In fact, according to several studies, large shares of servers run at average utilization levels below 20% [8, 9]. Resulting optimization potential must be explored, considering actual service demands and server capacities. Such data represent crucial input to consolidation decisions in order to identify orthogonal workload profiles over time and, lastly, to decide which services to place on which servers. As the solution space rapidly grows with the number of services, the problem becomes an intractable, NP-hard problem [1] and, therefore, error-prone if processed manually. Instead, it may be formulated as a combinatorial optimization problem with the objective to minimize the number of servers while avoiding overloads at any time. As a result, a dynamic and multi-dimensional bin packing problem can be formulated with

- $T \in \mathbb{N}_+$ the number of intervals which form the time dimension
- $S \in \mathbb{N}_+$ the number of servers with $0 < c_s$ and $0 < m_s$ defining the CPU and memory capacities of the servers for $1 \leq s \leq S$
- $V \in \mathbb{N}_+$ the number of services with a computing demand vector $\vec{c_v} = (c_{v,1}, \dots, c_{v,T})$, $0 \leq c_{v,t}$, $1 \leq t \leq T$ and a memory demand vector $\vec{m_v} = (m_{v,1}, \dots, m_{v,T})$, $0 \leq m_{v,t}$, $1 \leq t \leq T$, $1 \leq v \leq V$
- $X = (x_{s,v})_{1 \leq s \leq S, 1 \leq v \leq V}$ a binary matrix indicating if service v is allocated to server s
- $f(X) = \sum_{s=1}^{S} \text{sign}(\sum_{v=1}^{V} x_{s,v})$ the objective function to find an allocation matrix X that minimizes the number of required servers

Figure 1 exemplifies the allocation of services to servers. According to the objective function, the signum function returns 1 if a server was used at least one time while it returns 0 for unused servers. Therefore, the sum of used servers is to be minimized. However, a number of constraints apply to ensure quality of service:

- $\forall\ 1 \leq v \leq V$: $\sum_{s=1}^{S} x_{s,v} = 1$ a service must be allocated once, therefore, it's column in Figure 1 must sum up to 1
- $\forall\ 1 \leq t \leq T,\ 1 \leq s \leq S$: $\sum_{v=1}^{V} x_{s,v}\, c_{v,t} \leq c_s$ the sum of computing demands must not exceed the server's computing capacity at any time
- $\forall\ 1 \leq t \leq T,\ 1 \leq s \leq S$: $\sum_{v=1}^{V} x_{s,v}\, m_{v,t} \leq m_s$ the sum of memory demands must not exceed the server's memory capacity at any time

services (columns) / servers (rows)

$x_{i,j}$	1	2	3
A	1	0	1
B	0	0	0
C	0	1	0

Figure 1: Binary allocation of services to servers. Each service must be allocated to one server while a server may handle none, one, or multiple services. It is the objective to identify solutions which minimize the number of used servers. Overloads are avoided by considering dynamic workload profiles of services and maximum computing and memory capacities of servers.

The Digital Annealer: A Quantum-Inspired Computing Technology

Quantum computers promise to calculate problems that are probably never solvable with conventional CPU- and GPU-based technology. However, the use of quantum computers in real business scenarios is not economically viable at this time. With its computational architecture based on a digital circuit design, inspired by quantum phenomena, the Digital Annealer overcomes the current gap in quantum technology and paves the way for a much faster and more efficient approach of solving today's business problems.

Digital Annealer Business Applications

The Digital Annealer and the Digital Annealer Optimization Services are designed to solve large-scale combinatorial optimization problems.

There are numerous Digital Annealer business applications like:

- Route Optimization: Find the shortest path to a destination. Ideal for identifying the quickest transportation routes, public transportation planning, also applicable in factory robot route optimization.
- Industrial Equipment & Networks: Identify network design problems and develop optimal communication networks and oil & gas pipelines.
- Deliveries & Scheduling: Optimize delivery plans and scheduling. Ideal for postal and logistics services requiring delivery scheduling improvements, employee work hours scheduling and event coordination.
- Distribution & Warehouse Operations: Improve placement of in-stock parts in factories, stores and distribution warehouses.
- Disaster Recovery: Providing an efficient and quick recovery, rescue and aid plan in an event of calamities.

Digital Annealer Optimization Services

The Digital Annealer Optimization Services implement the Ising model in the fully connected form with high precision. Programming such a systems, the business problem has to be formulated as an Ising model or equivalently as a quadratic unconstrained binary optimization problem (QUBO). For more details and methods please refer to [1]. Optimum solution configuration will be associated with minimum energy of the QUBO. Providing quantum computing-like capabilities, the Digital Annealer can be operated at normal room temperature and small form factor. Its fully coupled 8,192-bit connectivity allows large-scale problem solving. 64-bit (2E64) gradations allow high accuracy in expressing combinatorial optimization problems.

In this case, optimization capabilities of the Digital Annealer are leveraged to optimize data centers. This is a multi-knapsack and multi resources problem where the Digital Annealer is allocating all services to available servers without any resource overflow and aiming to minimize the number of servers. Therefore, the Digital Annealer is ideal to solve such optimization problems.

Introduction of Solution Algorithms

In order to compare the Digital Annealer's performance in solving the server consolidation problem, state-of-the-art optimization algorithms for this problem are introduced and the problem-specific Ising formulation for the Digital Annealer is explained.

Genetic Algorithms Approach

The Digital Annealer is to be compared with two pure meta-heuristic approaches as well as one hybrid approach that combines meta-heuristics with problem-specific heuristics. Both meta-heuristics are genetic algorithms that mimic natural evolution of DNA on earth. For that reason, a set of solution candidates is randomly initialized, and altered in each generation using local (mutation) and global (recombination) search operators. A selection operator is used to choose the most promising candidates for the next generation. The selection depends on the fitness function that models the quality of a solution candidate. While both genetic algorithms used in the experiments rely on the same fitness function, they differ in the solution encoding and, thus, the genetic operators.

The first genetic algorithm, labeled GA, has been used e.g. by [2] and [3] for solving server consolidation problems. A solution is encoded in an integer vector $\boldsymbol{y}$ of length V such that any element $y_v = s,\ iff\ x_{s,v} = 1$ (i.e. for each service index a server index indicates to which server the service is to be allocated). While solution candidates are initialized by choosing random server indices, mutation of a single solution is implemented as a swap of server indices between two service indices. For recombination of two "parent" candidates, a server index is taken randomly from one parent for each service index (uniform crossover). The selection is performed in a tournament setting, in which a number of solutions are compared to each other and are selected for the next generation with a probability proportional to the solutions' relative fitness values. The GA was parameterized to use comma selection. Therefore, in contrast to plus selection, the offspring fully replaces the parent population. However, to keep the candidate with the highest fitness value so far, for the next generation, so called 1-elitism is used.

The second genetic algorithm, labeled GGA from grouping genetic algorithm, encodes a solution as a sequence of n sets for each of the servers (cf. e.g. [4, 5]). Thus, a set at index s contains all indices of services allocated to this server: $set_s = \{v | x_{s,v} = 1\}$. Candidates are initialized by assigning all service indices to a random set number. While for the mutation operator a random service index is shifted to another server set, recombination is a three-phased operator: First, the service sets are produced by performing uniform crossover between two parent candidates. Second, service indices that now appear twice in the solution are eliminated. Third, missing service indices are injected using a first-fit heuristic (services are allocated to the first set in which sufficient capacity remains). Again, comma-selection is used with 1-elitism, however, the technique of stochastic universal sampling is used to randomly select candidates for the next generation.

The hybrid algorithm, labeled GA_BF, is a combination of a genetic algorithm and a best-fit heuristic (cf. e.g. [6, 3]). In this approach, solution candidates are encoded as a permutation of service indices that represent the sequence in which a best-fit heuristic allocates these services. Thus, the fitness evaluation of a solution candidate is done by performing the allocation by the heuristic and evaluate the fitness of the created allocation similar to the other genetic algorithms. These candidates are initialized by creating random permutations and are mutated by swapping two service indices. For recombination, the edge crossover algorithm described by [7] is used to only create valid permutations, and tournament selection is used.

QUBO formulation for the Digital Annealer

The following constraints are formulated in a QUBO format.

1) Each service must be assigned to exactly one of the servers:

$$H1 = \sum_{v=1}^{V} \left(1 - \sum_{s=1}^{S} x_{s,v}\right)^2$$

2) CPU resources needed for services should not overflow on any server at all the times:
   $c_{v,t}$ is the CPU resource needed by service v at time t.
   c_s is the total CPU capacity of server s.
   $Y_{1,s,t}$ is CPU capacity not used in %, formulated with help of a so called slack variable.

$$H2 = \sum_{s=1}^{S} \sum_{t=1}^{T} \left(\sum_{v=1}^{V} \frac{x_{s,v} c_{v,t}}{c_s} + Y_{1,s,t} - 100\right)^2$$

3) Memory resources needed for services should not overflow on any server at all the times:
   $m_{v,t}$ is the memory resource needed by service v at time t.
   m_s is the total memory capacity of server s.
   $Y_{2,s,t}$ is the memory capacity not used in %.

$$H3 = \sum_{s=1}^{S} \sum_{t=1}^{T} \left(\sum_{v=1}^{V} \frac{x_{s,v} m_{v,t}}{m_s} + Y_{2,s,t} - 100\right)^2$$

4) z_s is 1 if any service is allocated to the server s.
   z_s is 0 if no service is allocated to the server s, as enforced by constraint.

$$H4 = \sum_{s=1}^{S} \left((1 - z_s) \sum_{v=1}^{V} x_{s,v}\right)$$

5) Solution with minimum servers is chosen:

$$H5 = \sum_{s=1}^{S} z_s$$

QUBO (*H*) will be the sum of all six constraints. A penalty factor is multiplied with each of the constraints. In case a constraint is not fulfilled, the corresponding energy is added as a penalty. Penalty A is very high as H1 represents services allocation to servers and must be fulfilled. Penalty B is associated with not allowing any resources overflow on servers. Penalty C is high as H4 is controlling the bits used for optimization constraint H5. Penalty D is very high as H5 represents the optimization of the problem. Finding the right balance between all the penalties is very important as the unbalanced penalties may adversely affect optimization or any QUBO constraint.

$$H = A \cdot H1 + B \cdot (H2 + H3) + C \cdot H4 + D \cdot H5$$

The Digital Annealer tries to find the configuration with the global minimum energy of the QUBO. The latter represents the optimized answer for the given problem. The Fujitsu's Digital Annealer cloud service has been used for this exercise to optimize the QUBO. From the optimal configuration returned by the Digital Annealer, the solution of the problem is formulated. Optimization with the Digital Annealer is also available as on premise service.

Aside from the constraints formulated above, there are other and more complex operational constraints which can be added to the problem. The Digital Annealer can very well optimize problems if further constraints are added to existing problems. The increase in complexity and size of the problem will barely affect the execution time and result quality unlike other algorithms.

Results of a Multi-Case Data Center Capacity Management Study

In the following, the presented solution algorithms are applied to solve server consolidation problems that are modeled from real-world cases. For that purpose, 246 data centers hosting enterprise systems have been monitored over a period of seven to 65 days. This information is used to define peak workload demands (computing and memory) for an hour-of-the-day season for each enterprise system. In combination with server capacities, a server consolidation problem is defined for each case.

The number of servers in the 246 cases differs from two to ten while between four and fifteen services have to be allocated. While the simplest case allows for 128 different allocations, the solution space comprises $1 * 10^{10}$ solution candidates for the most complex case.

In order to compare the quality of the solution algorithms, a solution quality index is defined that rewards smaller numbers of servers required in a solution and penalizes resource overflows. For that reason, the achieved number of required servers is increased by the value of five if computing overflows occur and again by five if memory overflows occur. On that basis, for each case the best quality index is computed, so that to each algorithm and case, the relative quality index difference can be assigned. The boxplot of the relative quality index difference of each algorithm is displayed in Figure 2:

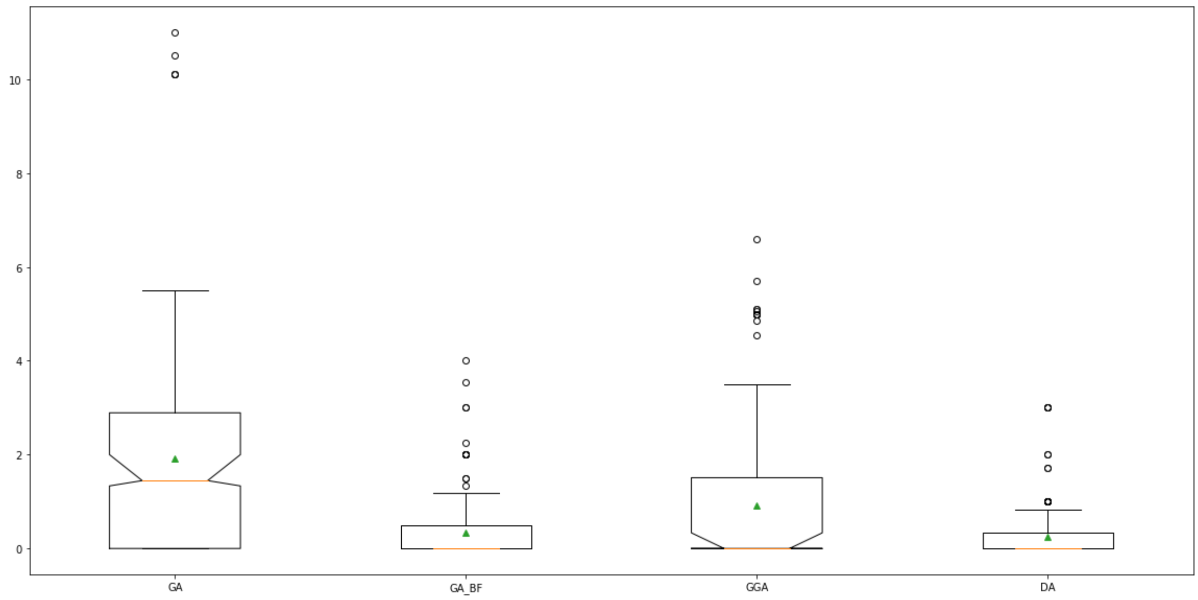

Figure 2: As a relative quality index difference of zero indicates that the best solution has been identified by this algorithm, it can be seen that for GA_BF, GGA, and DA the best solution is identified in more than half of the cases (indicated by the median - orange line), while the GA is considerably worse in this criterion. Average relative quality index difference of GA_BF and DA (green triangles) are significantly better than for GGA. However, while the DA seems to be a bit more stable than the GA_BF, a significant difference in average behavior cannot be concluded. This statement is also supported by pairwise comparisons of algorithmic performance, which are presented in Table 1.

S-values	DA	GA	GA_BF	GGA
DA		0.70731707	0.2398374	0.43089431
GA	0.11382114		0.09756098	0.3495935
GA_BF	0.14634146	0.69918699		0.40243902
GGA	0.23170732	0.41463415	0.23577236	

Y-values	DA	GA	GA_BF	GGA
DA		-0.31791694	-0.00538941	-0.06424894
GA	1.50016292		1.35864893	1.0845967
GA_BF	0.10013066	-0.29233895		-0.04355417
GGA	0.71257743	-0.09136984	0.58750779	

Table 1: In this context, the S-value refers to the ratio of cases in which an algorithm achieved a better solution quality than another algorithm. The Y-value is defined as the mean relative difference in solution quality between two algorithms [3]. Thus, the DA achieved better solutions than the GA_BF in nearly 24% of the cases, while GA_BF dominated DA only in 15% of the cases. However, DA achieved only a 0.5% better quality index than GA_BF on average. This supports the statement that DA and GA_BF have comparable performance in terms of solution quality with slight indicators that DA performs a bit better.

Nonetheless, referring to runtimes, the DA achieves this comparable solution quality in much less time. This can be seen in Figure 3, in which the box plots of algorithm runtimes in milliseconds are compared.

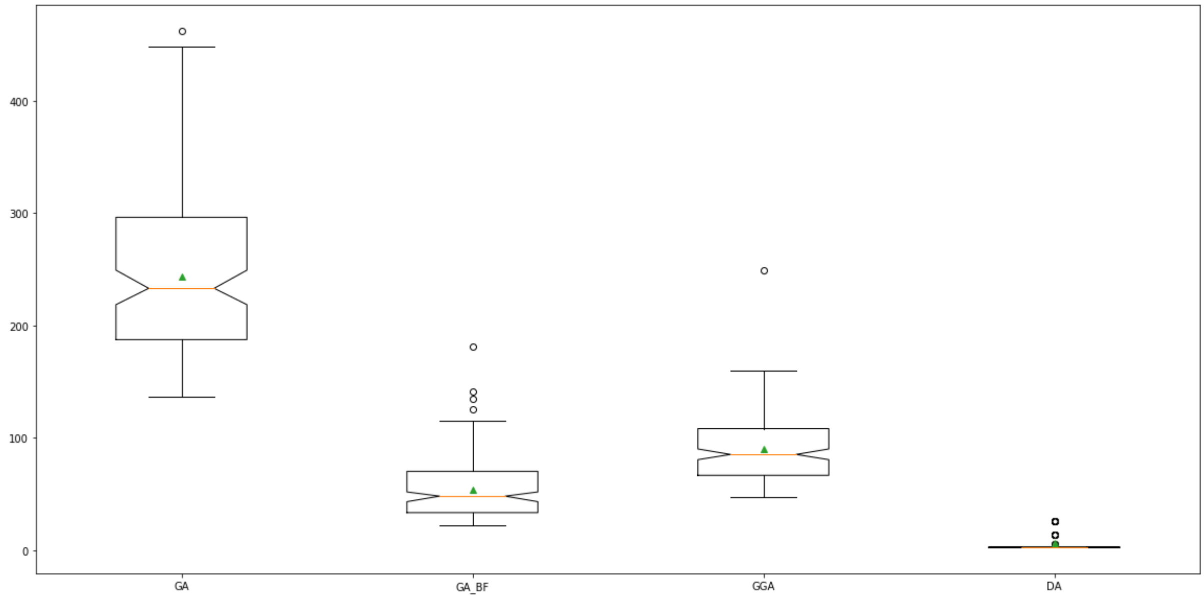

Figure 3: Comparing the DA performance with the best results from all other algorithms, the DA achieved state-of-the-art performance in 133 cases (54%) and identified better solutions in 33 cases (13%). Hence, the potential of the Digital Annealer to improve the server consolidation efforts in data centers could be revealed.

Conclusion

Fujitsu's Digital Annealer achieves state-of-the-art quality with high reliability. In 246 experiments from the domain of IT capacity management using real-world data, the Digital Annealer revealed optimal solutions in 67% of the cases, hence, more often than any other algorithm. In 33 of 246 cases (13.4%), the Digital Annealer has found a better solution than any other algorithm in any iteration, leading to the best average solution quality among the tested metaheuristics. Remarkably, the Digital Annealer identified feasible solutions faster by orders of magnitude when compared to state-of-the-art algorithms. In fact, the Annealer's runtime for the most complex case was lower than the runtime of any other tested metaheuristic for the least complex case. This makes the Digital Annealer an excellent technique to complement the solution portfolio of infrastructure optimization services for enterprise applications such as the Fujitsu SAP SystemInspection.

Executive Summary

Together with the Magdeburg Research and Competence Cluster for Very Large Business Applications, Fujitsu has conducted a multi-case study in order to challenge its new quantum-inspired optimization technology, termed **Digital Annealer**. The Digital Annealer is a cutting-edge optimization solution and, therefore, suitable to optimize modern data centers with respect to operational costs. As infrastructure optimization challenges of this kind are addressed by the Fujitsu SAP SystemInspection Service, the Digital Annealer was applied across 246 server consolidation use cases in order to find optimal SAP landscape designs and, therefore, to identify saving potential. When comparing the results to other state-of-the-art optimization algorithms, the Digital Annealer finds optima more often than any other algorithm and, in 13.4% of the cases, a better solution than any other algorithm. In the majority of the tested cases, the Digital Annealer outperforms state-of-the-art algorithms in terms of solution quality and runtime.

The **SAP SystemInspection Service** offers an efficient analysis and comprehensive consultation package for existing infrastructure environments. The provided analysis of current workloads, performance bottlenecks, and related resource consumption creates an excellent data basis for the optimization of SAP landscapes. As this challenge typically opens up vast solution spaces of NP-hard optimization problems, the quantum-inspired Digital Annealer represents a convenient technique to identify optimal or near-optimal designs. Such reorganizations of services deliver a positive contribution to the energy consumption of organizations and underline the Fujitsu Climate and Energy Vision.

Acknowledgement

Experiments performed for this white paper were conducted by the Magdeburg Research and Competence Cluster (MRCC) for Very Large Business Applications (VLBA) at the Otto-von-Guericke-University Magdeburg and the Digital Incubation team, Fujitsu Technology Solution, Germany. Special thanks to Dr. Fritz Schinkel, Andreas Rohnfelder for giving us the opportunity and support to solve this problem using the Digital Annealer. We are thankful to Markus Kirsch for his technical reviews and Alexander Kaffenberger for initiating this project. Also, we would like to thank Marco Edel and Jürgen Ellwanger for the ongoing efforts to create and manage a scientific environment which enabled to study cases and conduct experiments for this white paper. Finally, we thank the members of the Chair of Business Informatics at the Otto-von-Guericke-University Magdeburg for the continuous support of applied research in collaboration with Fujitsu.

Authors

Sushant Uttam Dhamnekar[1], Dr. Sascha Bosse[2], Dr. Hendrik Müller[2], Dr. Christian Münch[2], Prof. Klaus Turowski[2]

References

1) Lucas, Andrew, 2014, Ising formulations of many NP problems, Frontiers in Physics Feb. 2014, https://doi.org/10.3389/fphy.2014.00005
2) Rolia, J., Andrzejak, A., and Arlitt, M. (2003). Automating enterprise application placement in resource utilities. In Brunner, M. and Keller, A. (eds.), 14th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM), pages 118–129.
3) Stillwell, M., Schanzenbach, D., Vivien, F., and Casanova, H. (2010). Resource allocation algorithms for virtualized service hosting platforms. Journal of Parallel and Distributed Computing, 70:962–974.
4) Xu, J. and Fortes, J. A. B. (2010). Multi-objective virtual machine placement in virtualized data center environments. In IEEE/ACM International Conferenc on Cyber, Physical and Social Computing (CPSCom), Hangzhou, China.
5) Falkenauer, E. and Delchambre, A. (1992). A genetic algorithm for bin packing and line balancing. In Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on, pages 1186–1192.
6) Liu, D. S., Tan, K. C., Huang, S. Y., Goh, C. K., and Ho, W. K. (2008). On solving multiobjective bin packing problems using evolutionary particle swarm optimization. European Journal of Operational Research, 190:357–382.
7) Larranaga, P., Kuijpers, C. M. H., Murga, R. H., Inza, I., and Dizdarevic, S. (1999). Genetic algorithms for the travelling salesman problem: A review of representations and operators. Artificial Intelligence Review, 13(2):129–170. (Cited on page 93.)
8) Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., and Lintner, W. (2016). United states data center energy usage report. URL https://escholarship.org/uc/item/84p772fc
9) Koomey, J. and Taylor, J. (2017). Zombie/comatose servers redux. Technical report, Stanford University, Anthesis, Koomey Analytics.

[1] Fujitsu Technology Solutions
[2] Magdeburg Research and Competence Cluster for Very Large Business Applications at the Otto von Guericke University Magdeburg

Fujitsu SAP SystemInspection Service

Overview

Combinatorial Optimization Services for Modern Data Centers are available for organizations as SAP SystemInspection Service.

The SAP SystemInspection Service is one of our offerings with a strong focus on SAP Technology such as SAP NetWeaver and SAP HANA. It offers an efficient analysis and comprehensive consultation package for existing infrastructure environments at a fixed price. The goal is to get a complete understanding of the current workload, performance and related resources consumption & distribution in order to provide measures and guidance to optimize your SAP Landscapes according to the business strategy & requirements.

The SAP SystemInspection Service analyzes the customer's operational SAP IT Infrastructure, builds an accurate mapping of the actual load profile and formulates an inventory list of the entire SAP IT infrastructure. It creates a holistic view and transparency of the performance and utilization of individual IT components and associated SAP systems. The findings are presented and discussed during the consultation session. Recommendations or Instructions for actions are provided. The SAP SystemInspection Service is supported by heuristics and metaheuristics algorithms that were developed by the University of Magdeburg specifically for Fujitsu, addressing the Dynamic Priority-based Workload Consolidation Problem (DPWCP).

Benefits

- Avoid over-provisioning and unnecessary investments in equipment
- Increase the quality of services by helping to eliminate or circumvent performance bottlenecks
- Delivers individual and short-term results with minimum effort for the customer
- Insightful base for capacity planning of upcoming new requirements and to make strategic decisions
- Gives clear recommendations for optimizing the IT infrastructure for SAP

Contact
FUJITSU
Fujitsu Technology Solutions GmbH
Mies-van-der-Rohe-Strasse 8, 80807 München, Deutschland
E-Mail: cic@ts.fujitsu.com
Website: www.fujitsu.com/digitalannealer