MemorAI: Energy-Efficient Last-Level Cache Memory Optimization for Virtualized RANs

Ethan Sanchez Hidalgo[†], J. Xavier Salvat Lozano^{*}, Jose A. Ayala-Romero^{*},

Andres Garcia-Saavedra*, Xi Li * and Xavier Costa-Perez*^{†‡}

NEC Laboratories Europe*, i2CAT Foundation[†], ICREA Foundation[‡]

Email: ethan.sanchez@i2cat.net, {josep.xavier.salvat, jose.ayala, andres.garcia.saavedra, xi.li, xavier.costa}@neclab.eu

Abstract-The virtualization of Radio Access Networks (vRAN) is well on its way to become a reality, driven by its advantages such as flexibility and cost-effectiveness. However, virtualization comes at a high price - virtual Base Stations (vBSs) sharing the same computing platform incur a significant computing overhead due to in extremis consumption of shared cache memory resources. Consequently, vRAN suffers from increased energy consumption, which fuels the already high operational costs in 5G networks. This paper investigates cache memory allocation mechanisms' effectiveness in reducing total energy consumption. Using an experimental vRAN platform, we profile the energy consumption and CPU utilization of vBS as a function of the network state (e.g., traffic demand, modulation scheme). Then, we address the high dimensionality of the problem by decomposing it per vBS, which is possible thanks to the Last-Level Cache (LLC) isolation implemented in our system. Based on this, we train a vBS digital twin, which allows us to train offline a classifier, avoiding the performance degradation of the system during training. Our results show that our approach performs very closely to an offline optimal oracle, outperforming standard approaches used in today's deployments.

Index Terms—RAN virtualization, Last-Level Cache Memory, Noisy Neighbour Problem, Digital Twin

I. INTRODUCTION

RAN virtualization is a crucial technology for reducing the Total Cost of Ownership (TCO) of 5G RAN infrastructure [1]-[3]. Virtualized RANs (vRANs) are expected to import the advantages of network function virtualization (NFV), such as exploiting general-purpose computing platforms and shortening deployment cycles. Collaborative efforts like the carrier-led O-RAN alliance have stimulated the market and the research community to develop innovative solutions that incorporate the adaptability and cost-effectiveness of NFV right into the edge of mobile networks [4]. Nonetheless, while shared computing platforms offer enhanced flexibility and cost-effectiveness, they also introduce challenges for 5G base stations, as they compromise the predictability offered by dedicated platforms [5]–[7]. The term *noisy neighbor problem* has been coined to refer to the issue when shared resources are consumed in extremis, meaning that another function restricts one virtualized function's resources. This problem has motivated substantial research over the years [6], [8]-[10].







Fig. 2: Energy consumption as a function of the computing load.

Virtual RANs are not alien to the *noisy neighbors problem* as shown in our previous work [11]. When deploying multiple virtual Base Stations (vBSs) instances in a vRAN system, the usage of computing resources increases. We confirm this by deploying multiple vBS instances using Docker containers into a pool of computing cores in a shared off-the-shelf server. In contrast to [11], in these experiments we configured the pool of computing cores to retain as much predictability as possible (see §III for more details).

Fig. 1 depicts the aggregated computing usage per core in our vRAN platform as a function of the number of deployed vBS under maximum traffic load in uplink (UL) and downlink (DL). The bars in yellow show the expected usage assuming perfect resource isolation in place. We compute these by linearly scaling up the CPU usage of a single vBS instance. The red bars show the actual CPU consumption, which unveils an exponentially-growing overhead induced by the aforementioned *noisy neighbors problem*.

As we explained in [11], the increased computing overhead has its roots in the lack of cache memory isolation. The computing overhead poses an issue of the increase of the total energy consumption of the vRAN platform. Fig. 2 shows the relationship between the normalized energy consumption on top of the system's baseline (i.e. idle) consumption of our vRAN platform as a function of the total computing load. The computing load and the energy are linearly related [12], [13]. Therefore, it is key to minimize the computing usage of our vRAN platform to keep operational energy costs low.

In our previous work [11], we developed a solution to dimension the computing capacity of a vRAN system, considering the increased computing usage due to the *noisy neighbors problem*. However, this work does not consider minimizing the increased overhead but adapting to it. In this paper, we seek to reduce the computing overhead as much as possible.

The work was supported by the European Commission through Grant No. SNS-JU-101097083 (BeGREEN). This work was also supported by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union – NextGeneration EU, in the framework of the Recovery Plan, Transformation and Resilience (PRTR) (Call UNICO I+D 5G 2021, ref. number TSI-063000-2021-3) Additionally, it has been supported by MINECO/NG EU (No. TSI-063000-2021-7) and the CERCA Programme.

Memory type	Access latency [15]-[17]	Size per core
L1 cache	4-6 cycles	64 KB
L2 cache	14 cycles	256 KB
L3 (LLC) cache	50-70 cycles	2 MB
RAM	~ 120 - 600 cycles	2-4 GB

TABLE I: Access and cache miss latency.

We begin studying how isolating the different cache memory levels influences the total system computing consumption. We found that vBSs traffic demand is virtually orthogonal to using cache resources, i.e., vBSs use as much cache memory as possible. However, the utility of cache memory is different for different traffic demands. The vBSs with higher demands and Signal-to-Noise-Ratio (SNR) can reduce their computing usage more when they have more cache memory available. Thus, developing a solution that strategically allocates the cache memory resources to the different vBS instances according to their demands is key to minimize energy consumption. Due to the complex relationship between computing and radio resources [14], we propose a novel approach that can effectively allocate the cache resources to minimize the total computing usage and consequently reduce the energy consumption of vRAN platforms.

II. BACKGROUND: MEMORY IN GENERAL-PURPOSE COMPUTING PLATFORMS

Cache memory [15] bridges the speed gap between RAM and the CPU itself and it is organized into various levels based on speed and size. The first two cache memory levels called L1 and L2 cache are the closest and fastest to the system, although its capacity is limited. Each physical core has its dedicated L1 and L2 cache. L1 is faster than L2 but lower in size. Lastly, the L3 cache, also known as the *Last Level Cache (LLC)*, is the slowest CPU's cache and is shared among all computing cores [15]–[17]. Table I shows the access latency and the available size per core of different cache levels in the Intel Skylake architecture.

Cache memory is organized into several associative sets also known as cache ways. A cache way is a set of lines or blocks that can store data from a specific location in the main memory. Tools such as Intel CAT¹ enables allocating LLC cache ways to specific cores or processes.

In a general-purpose computing platform, when a core executing a thread requires frequently used memory blocks, it loads them into the cache for quicker access. Consequently, if a thread references a memory block that is not present in the cache, the core initiates an interrupt known as a "cache miss" and searches for the data in a higher memory level. This incurs additional CPU cycles which increase the total execution time of a computing thread.

III. EXPERIMENTAL ANALYSIS

A. Cache memory isolation

To evaluate the impact of cache memory resources on the *noisy neighbor problem* in vRANs' energy consumption, we measure the computing usage and low-level cache metrics



Fig. 3: vRAN testbed.

deploying a different number of vBS instances in different scenarios in our vRAN platform. Fig. 3 shows our experimental vRAN platform. Therein, we deploy the different vBS instances in an isolated pool of computing cores from an Intel Xenon E5-2650 v4 CPU @ 2.20GHz in a shared off-the-shelf server. Each vBS has its dedicated RF radio head connected to one UE, which is used to emulate the aggregated cell load. We configured the pool of computing cores to retain as much predictability as possible: (i) we isolated 12 cores with 12 dedicated Last Level Cache (LLC) ways, (ii) the system can only use C0/C1 C-states and we turned off hardware P-states and (iii) we deactivated Hyper-threading. We then initiate bidirectional data flows, both uplink (UL) and downlink (DL), with maximum load and good wireless channel conditions between each vBS instance and different user equipment (UE). We consider the following scenarios:

- *Ideal*. We compute the CPU usage scaling linearly the usage of a single vBS instance assuming that the cache memory size also scales linearly.
- No isolation. We deploy an increasing quantity of vBS instances without any cache memory isolation mechanism.
- *Pinning.* L1 and L2 cache levels are dedicated per core. We pin different vBS instances to distinct cores to assess the impact on L1 and L2 cache isolation. To facilitate comparisons, the number of cores assigned to each vBS is given by:

cores per vBS =
$$\left| \frac{\text{total cores}}{\text{deployed vBSs}} \right|$$
 (1)

where the total number of cores equals 12 and the number of deployed vBS increases from 1 to 5. Note that when 5 vBSs are deployed, each vBS is assigned to 2 cores and two free cores are left.

• *Pinning* + *LLC isolation*. We perform the L3 cache allocation in the same manner as the CPU pinning. We allocate the total number of cache ways equally for every single vBS, i.e. the number of cache ways per vBS is given by:

cache ways per vBS =
$$\left\lfloor \frac{\text{total cache ways}}{\text{deployed vBSs}} \right\rfloor$$
 (2)

In our platform, there are a total of 12 cache ways. We use Intel CAT to allocate the LLC cache ways to the corresponding computing cores.

To measure the computing usage we use the /proc filesystem to read and store periodically the computing time for the different threads of a vBS on the computing cores allocated. On the other hand, we used the tool perf to measure the

¹https://github.com/intel/intel-cmt-cat



Fig. 4: Comparison of the aggregated per-core usage with # of vBS instances showing the "No isolation", the "Pinning" and the "Pinning + LLC isolation" scenarios.





number of instructions per cycle (IPC) and the number of cache misses per 1k instructions (MPKI) by one vBS.

Fig. 4 shows the measured computing usage in the different scenarios. We observe that, for the no isolation configuration, the computing usage increases by approximately 50% compared to the ideal case, increasing the energy consumption. This also impacts the IPC and MPKI, as shown in Fig. 5-6, respectively. We observe that the IPC decreases and the cache misses increase as more instances are deployed. There is a sixfold increase in the cache misses when transitioning from 1 to 5 vBSs. This increase also impacts the number of cycles required to execute the same number of instructions, decreasing the IPC.

The *Pinning* configuration shows a lower computing usage than the *No isolation* but still higher than the ideal case. The computing usage decrease is significant considering that L1 cache ways are approximately 100 times lower in size than L3 cache ways and L2 cache ways are 10 times lower in size than L3 cache ways. In Fig. 5 we can observe a higher IPC across any number of deployed vBSs compared to the *No isolation* scenario. However, Fig. 6 shows the same number of cache misses for all the cases. This might seem shocking at first glance, but as we have previously detailed L1 and L2 cache way size is very low compared to L3.

Finally, to study the effect of the L3 cache isolation. In Fig. 4, the *Pinning* + *LLC isolation* configuration does not improve the computing usage. Specifically, the computing usage is the same for all the cases except for the case with 5 vBS, in which it is slightly higher. The reason behind this behavior is that we only allocate 10 out of 12 L3 cache ways to all vBSs for the case of 5 vBSs, while in the *Pinning* configuration, all vBSs have all the L3 cache memory available, using on average 2.4 L3 cache ways.

B. LLC occupancy and utility

We now study how the allocation of LLC cache ways impacts the computing usage of a vBS instance. First, we



Fig. 7: LLC occupancy in % as a function of the total demand for different SNR cases in UL and DL.



Fig. 8: Computing usage as a function of the LLC allocated cache ways for different SNR and in UL and DL.

measure the percentage of LLC cache memory used of the total LLC memory allocated to a vBS as a function of the traffic demand. Fig. 7 shows the LLC occupancy with 12 and 2 LLC cache ways for different SNR values and traffic demands in uplink and downlink. The *high* series depict the LLC occupancy with a high SNR environment while the *low* series depict the LLC occupancy with a low SNR environment. We can see that the total LLC occupancy is above 80% for uplink and downlink. Also, the LLC occupancy is almost orthogonal to the demand of the vBS, yielding an increase of a 2 - 6% when the demand goes from 10% to 100%.

Finally, Fig. 8 depicts the computing usage of a vBS as a function of the L3 cache ways when we deploy it using 3 cores. Fig. 8 depicts the computing usage for the different Modulation Coding Schemes (MCSs) with a low traffic demand (i.e. 20% of the total demand) and high traffic demand (i.e. 100% of the total demand) for uplink and downlink. We observe that there are significant differences between the achievable computing usage reduction. For high traffic both in uplink and downlink, we can achieve a more significant computing usage reduction. On the contrary, a vBS that processes a low traffic demand can attain lower gains in terms of computing usage. We conclude that LLC resources have different *utility* depending on the vBS context.

This contrasts with Fig. 7 which shows that the vBS makes full use of the cache memory regardless of the demand. Thus, if there is no LLC cache allocation mechanism, all vBSs deployed will be using the same amount of LLC cache memory on average. It is key to strategically distribute the



Fig. 9: Computing usage and decoding time of a vBS with max. UL and DL load with mild MCS over different SNR conditions.

LLC cache ways among the vBSs deployed boosting its utility to minimize the computing usage and therefore the energy consumption depending on the traffic demands.

C. The problem

Quantifying the LLC cache memory utility is a challenging task. The LLC cache utility depends on the computing demands of the vBS which are influenced by various factors, including traffic demand in both the downlink (DL) and uplink (UL), the signal-to-noise ratio (SNR) of each wireless link, and the specific Modulation Coding Scheme (MCS) utilized for communication. All these elements interact in a complex manner [14], [6]. Fig. 9a depicts the relative mean core usage of the vBS, and shows that, given a MCS, lower SNR regimes demand a higher amount of computing resources. The underlying reason is the iterative nature of the forward-errorcorrection (FEC) algorithms - signals received with lower SNR require a higher number of FEC iterations to decode the transported codeword successfully. This is confirmed by Fig. 9b, which shows the amount of time taken by the decoder to finish its task for every transport block.

IV. PROBLEM FORMULATION

We consider a vRAN platform comprising M_{cores} computing cores and $N_{\rm LLC}$ LLC cache ways. We consider that $N_{\rm vBS}$ vBS instances are deployed in the platform. Every vBS i in the vRAN platform has both a UL and DL traffic demand denoted by d_i^{UL} and d_i^{DL} , respectively; a SNR s_i ; and an MCS in UL and DL denoted by m_i^{UL} and m_i^{DL} , respectively, that the radio scheduler selects depending on s_i . The vBS *i* has a set of isolated cores P_i such that $|P_i| > 0$ and a number of dedicated LLC cache ways n_i^{LLC} , where $n_i^{LLC} \ge 1 \quad \forall i$. We need to satisfy that $\sum_i |P_i| \le M_{cores}$ and $\sum_i n_i^{LLC} = N_{LLC}$. We define c_i as the computing use of vBS i.

We define $\mathbf{x}_i := (d_i^{\text{UL}}, d_i^{\text{DL}}, s_i, m_i^{\text{UL}}, m_i^{\text{DL}})$ as the context of the vBS *i*. Moreover, we define f_i as the function that maps \mathbf{x}_i and n_i^{LLC} to the computing usage c_i .

Finally, we define vector \mathcal{X} which concatenates the context vectors from all vBS as $\mathcal{X} := (\mathbf{x}_1, \mathbf{x}_2, \dots \mathbf{x}_{N_{\text{vBS}}})$. Also, we define $\mathcal{P} := (P_1, \ldots, P_{N_{\text{vBS}}})$ and $\mathcal{N} := (n_1^{LLC}, \ldots, n_{N_{\text{vBS}}}^{LLC})$ as the vectors with the core set allocation and the LLC cache ways allocation on a vRAN platform. Similarly, we define the function f which maps $(\mathcal{X}, \mathcal{N})$ to the total computing usage $C^{\text{vRAN}} \in [0, M_{cores}]$ of the vRAN platform. We define



Fig. 10: MemorAI optimization framework.

the problem of optimizing the computing set and the LLC allocation as:

$$\begin{array}{ll} \min_{\mathcal{N}} & f(\mathcal{X}, \mathcal{N}) & (3) \\ \text{subject to} & \sum_{i} n_i^{LLC} = N_{LLC}. \end{array}$$

As explained in §I, the computing usage and the energy consumption are proportional. Thus, minimizing the computing usage function f also minimizes the total energy consumption. Note that, in our problem, the assignment of CPU cores to vBS \mathcal{P} is already given. We select a fixed value of \mathcal{P} based on our experimental insights and previous works on this topic [11], [14] (see Sec. VI for more details). Note that, in the problem (3), the optimal LLC cache allocation depends on the context X, whose dimensionality increases depending on the number of vBSs. Moreover, the optimal action is also dependent on the number of active vBS, which may change over time. To avoid the exploration burden of learning algorithms (e.g., reinforcement learning) that can lead to suboptimal configurations, we decompose the problem and use a digital twin system.

V. MEMORAI

To solve (3), we propose an optimization framework which we call MemorAI. This framework considers discrete decision intervals denoted by $t \in \{1, 2, \dots, T\}$. At the beginning of each decision interval, our solution receives $\mathcal{X}^{(t)}$, and decides the optimal LLC allocation $\mathcal{N}^{*(t)}$ across all vBSs, to minimize the energy consumption. MemorAI is composed of a set of Digital Twins (DTs) and a Neural Network (NN) classifier. Fig. 10 shows our optimization framework.

A. Digital Twin

r

Thanks to full vBS isolation via pinning it to dedicated cores and LLC cache ways allocation, we make the observation that $C^{\text{vRAN}} = \sum_{i} c_i$ as there are no joint effects between vBSs. This observation implies that $f(\mathcal{X}, \mathcal{N}) = \sum_{i} f_i(\mathbf{x}_i, n_i^{LLC}).$ As there is no interaction among vBS in terms of computing usage, we can create DTs of independent vBS. Thus, we can mirror their behavior in a safe and controlled environment for testing and learning. Each DT can model the particularities of each vBS (e.g., different implementations protocol stacks) and we can emulate the complex interactions of the full vRAN system. We create a DT using operational data from one vBSs. Using each vBS's DTs, we can lower the time cost to generate a dataset for a number of vBS as we aggregate the results of each DT. Note that without the DTs, the amount of data needed to create a training dataset increases exponentially with the number of vBS (curse of dimensionality).



To create the Digital Twin of one vBS we used our vRAN platform to generate a training dataset with a fixed computing core set |P|. Note that we select a set of cores such that vBS can correctly operate. This makes f_i a continuous function. Each dataset sample contains a 4-tuple with $(\mathbf{x}, P, n^{\text{LLC}}, c)$. Using this dataset, we built up a DT using a Neural Network (NN) which approximates c using $(\mathbf{x}, P, n^{\text{LLC}})$ i.e. it approximates f_i minimizing the Minimum Squared Error (MSE). We denote the DT function as \hat{f}_i^D . Fig. 10 (1), shows the DT training step.

B. NN Classifier

The DT allows us to evaluate the CPU usage of different configurations very accurately without having to use the real system. Therefore, we can perform an exhaustive search to find the optimal LLC allocation \mathcal{N}^* for a set of contexts $(\mathcal{X}, \mathcal{P})$. Note that the size of the set of possible LLC cache allocations is $|\mathcal{N}| = {N_{LLC} - 1 \choose N_{vBS} - 1}$. Fig. 10 (2), shows how using the different DTs we generate the previous data set.

Using this data set, we train a fully connected NN to predict \mathcal{N}^* for a given context $(\mathcal{X}, \mathcal{P})$. Specifically, we solve a multi-class classification problem using the cross-entropy loss function. We denote the classifier function as \hat{f}^C . Note that our solution is very flexible to changes in the system (e.g., upgrades in the implementation of the software stack, deployment of new vBSs, etc.). In those cases, after having the digital twin modeling, we can easily retrain the NN classifier offline without degrading the performance of the system. Finally, Fig. 10 (3), shows the classifier training step of our optimization framework.

VI. PERFORMANCE EVALUATION

In this section, we evaluate the performance and potential savings of our approach. We carry out the evaluation for a $N_{vBS} = 5$ vBS deployment. We used PyTorch² to implement the Digital Twin and the Classifier.

A. Training Evaluation

1) Digital Twin: We implemented the DT of one vBS using a NN with three hidden layers of sizes $\{256, 128, 64\}$ respectively with a ReLU activation function. Moreover, we stop the training iterations using an early stopping mechanism to prevent overfitting [18]. The early stopping mechanism stops the training after the value of the loss function in a validation data set has not improved during the last N training iterations. N is usually referred to as patience.

²https://pytorch.org/



Fig. 13: Energy savings compared to different benchmarks and different number of cache ways for a 15 min decision interval.

Fig. 11 shows the MSE loss value on the training and validation data sets of the DT as a function of the training iterations. We selected a patience of N = 10 and trained the model during 70 iterations.

2) NN Classifier: On the other hand, we implemented the Classifier as a NN with 4 hidden fully connected layers of sizes $\{512, 384, 384, 512\}$ respectively. Each layer also used a ReLU activation function and we introduced a 0.2 dropout probability during training. We also use the early stopping mechanism with N = 50 of patience.

Fig. 12 shows the cross entropy loss on the training and validation data set as a function of the number of training iterations. We train the model until iteration 300 (due to the early stopping mechanism). The training loss is higher than the validation loss due to the dropout layers. Also, the classifier achieves a 92.1% accuracy on our testing data set.

B. Performance Benchmark

We design *MemorAI* to operate in the Non-Real Time RIC of the O-RAN architecture as an rApp [19]. Based on this, we selected a time granularity of 15 minutes for our evaluation [20]. We generated a data set with random context data and compared *MemorAI* against the following approaches:

- **Random**: We select the allocation of cache ways for each vBS randomly;
- Equal partition: All the vBSs get allocated the same number of cache ways. Extra cache ways are left unallocated;
- Weighted: Each vBS gets allocated a number of cache ways as proportionally to its total demand as:

$$n_i^{LLC} = \frac{d_i^{\mathrm{UL}} + d_i^{\mathrm{DL}}}{\sum_j d_j^{\mathrm{UL}} + d_j^{\mathrm{DL}}} \cdot N_{\mathrm{LLC}}$$

Fig. 13a shows the energy savings in kilojoules (kJ) of our solution and the optimal strategy with respect to the different benchmarks in a decision interval of 15 minutes. Our solution outperforms the three benchmarks in terms of energy while producing almost the same results as the optimal solution. Thus, *MemorAI* understands better the utility of LLC cache partitioning than benchmark strategies. Our solution yields higher savings up to 1 kJ compared to the random strategy, up to 0.35 kJ compared to the equal partitioning, and up to 0.36 kJ compared to the weighted strategy shows a lower power consumption, it does not scale the LLC cache properly when there is an imbalance

between UL and DL demands. In these cases, our approach achieves the highest gains with respect to this strategy.

Finally, Fig. 13b shows the attained gains when the system has only 8 cache ways available. In that case, the random approach performs better because the number of configurations is lower and therefore the probability of selecting the optimal configuration is higher. The savings compared to the equal partition approach are higher as the benchmark can only allocate one cache way per vBS in this scenario. Compared to the weighted strategy, we still observe higher savings.

VII. RELATED WORK

The noisy neighbor problem has been studied for different types of virtualized resources. The works in [7], [21] propose different solutions to effectively isolate the networking stack. The authors in [21] show that containers' computing time processing packets via interrupts is not correctly accounted. Therefore, the authors develop a solution to called Iron which effectively charges the computing time used for processing packets to each container. Authors in [7] develop PicNIC a predictable virtualized NIC abstraction that effectively provides predictable performance to cloud providers.

The works in [8]–[10] develop different solutions to partition the memory resources to effectively isolate different tenants using the same infrastructure. The work in [8] develops a solution to estimate the slowdown of an application due to the interference from other applications and propose different strategies to share the memory resources. In [10] authors develop a clustering solution to fairly partition the LLC cache ways for different applications using Intel CAT. Moreover, in [9] the authors develop a solution to fairly allocate LLC cache and memory bandwidth for workload consolidation.

Finally, [6], [11], [22] tackle the *noisy neighbor problem* in the context of vRANs. In our previous work in [11] we showed that vRANs experience an increased computing consumption due to interference in the cache memory that can lead users to lose connectivity. We designed AIRIC, an AI controller that dimensions the computing capacity of a vRAN system considering the computing overhead. The authors in [6] enhance the physical layer pipeline of operations of a vBS so that it is suitable to run in non-deterministic computing platforms. Finally, in [22] authors develop a solution to increase the CPU utilization in vRAN platform opportunistically co-locating non-5G workloads while ensuring correct operation.

VIII. CONCLUSIONS

Cache memory is a key resource for vRANs to reduce energy consumption. Non-isolated access to cache memory resources increases energy consumption, making less attractive the advantages of virtualization. In our work, we have studied how the different mechanisms for cache memory isolation decrease the energy consumption of a vRAN platform. Then, we proposed *MemorAI* which strategically allocates LLC resources to minimize energy consumption. *MemorAI* comprises a digital twin and a neural network classifier, providing a very efficient and flexible solution. *MemorAI* achieves almost optimal performance and can attain significant energy savings when compared with other strategies.

REFERENCES

- M. Masoudi, S. S. Lisi, and C. Cavdar, "Cost-effective migration toward virtualized c-ran with scalable fronthaul design," *IEEE Systems Journal*, vol. 14, no. 4, pp. 5100–5110, 2020.
- [2] F. W. Murti *et al.*, "An optimal deployment framework for multicloud virtualized radio access networks," *IEEE Transactions on Wireless Communications*, vol. 20, no. 4, pp. 2251–2265, 2020.
- [3] Samsung, "Virtualized Radio Access Network: Architecture, Key technologies and Benefits." *Technical Report*, 2019, Link.
- [4] O-RAN Alliance, "Cloud Architecture and Deployment Scenarios for O-RAN Virtualized RAN (O-RAN.WG6.CADS-v04.00)," Technical Report, Oct. 2022.
- [5] A. Tootoonchian et al., "ResQ: Enabling SLOs in Network Function Virtualization," in Proceedings of the 15th USENIX NSDI, 2018, pp. 283–297.
- [6] G. Garcia-Aviles, A. Garcia-Saavedra, M. Gramaglia, X. Costa-Perez, P. Serrano, and A. Banchs, "Nuberu: Reliable ran virtualization in shared platforms," in *Proceedings of the 27th Annual International Conference* on Mobile Computing and Networking, 2021, pp. 749–761.
- [7] P. Kumar, N. Dukkipati, N. Lewis, Y. Cui, Y. Wang, C. Li, V. Valancius, J. Adriaens, S. Gribble, N. Foster *et al.*, "Picnic: predictable virtualized nic," in *Proceedings of the ACM Special Interest Group on Data Communication*, 2019, pp. 351–366.
- [8] L. Subramanian et al., "The application slowdown model: Quantifying and controlling the impact of inter-application interference at shared caches and main memory," in *Proceedings of the 48th International Symposium on Microarchitecture*, 2015, pp. 62–75.
- [9] J. Park, S. Park, and W. Baek, "Copart: Coordinated partitioning of last-level cache and memory bandwidth for fairness-aware workload consolidation on commodity servers," in *Proceedings of the Fourteenth EuroSys Conference 2019*, 2019, pp. 1–16.
- [10] V. Selfa *et al.*, "Application clustering policies to address system fairness with intel's cache allocation technology," in 2017 26th international conference on parallel architectures and compilation techniques (pact). IEEE, 2017, pp. 194–205.
- [11] J. X. Salvat Lozano, A. Garcia-Saavedra, X. Li, and X. Costa Perez, "AIRIC: Orchestration of virtualized radio access networks with noisy neighbours," *Accepted for publication in IEEE Journal on Selected Areas* in Communications, 2023.
- [12] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," ACM SIGARCH computer architecture news, vol. 35, no. 2, pp. 13–23, 2007.
- [13] C. Lefurgy et al., "Server-level power control," in Fourth International Conference on Autonomic Computing (ICAC'07). IEEE, 2007, pp. 4–4.
- [14] J. A. Ayala-Romero *et al.*, "vrain: Deep learning based orchestration for computing and radio resources in vrans," *IEEE Transactions on Mobile Computing*, vol. 21, no. 7, pp. 2652–2670, 2022.
- [15] B. Jacob, D. Wang, and S. Ng, *Memory systems: cache, DRAM, disk.* Morgan Kaufmann, 2010.
- [16] J. Patterson, "Modern microprocessors: A 90 minute guide!" Cortex, vol. 15, p. A57, 2003.
- [17] U. Drepper, "What every programmer should know about memory," *Red Hat, Inc*, vol. 11, p. 2007, 2007, Link.
- [18] L. Prechelt, "Early stopping-but when?" in *Neural Networks: Tricks of the trade*. Springer, 2002, pp. 55–69.
- [19] O-RAN Alliance, "O-RAN Non-RT RIC Architecture 3.0 (O-RAN.WG2.Non-RT-RIC-ARCH-R003-v03.00)," Technical Report, Jun. 2023.
- [20] C. Marquez et al., "How should i slice my network? a multi-service empirical evaluation of resource sharing efficiency," in Proceedings of the 24th Annual International Conference on Mobile Computing and Networking, 2018, pp. 191–206.
- [21] J. Khalid, E. Rozner, W. Felter, C. Xu, K. Rajamani, A. Ferreira, and A. Akella, "Iron: Isolating network-based {CPU} in container environments," in 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18), 2018, pp. 313–328.
- [22] X. Foukas and B. Radunovic, "Concordia: Teaching the 5g vran to share compute," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 580–596.