Demo: Smart Waste Disposal with Edge-Cloud Continuum Architecture

Josef Spillner
josef.spillner@zhaw.ch
Zurich University of Applied Sciences
Winterthur, Switzerland

Mengwei Liu
mengwei.liu@kaust.edu.sa
King Abdullah University of Science and Technology
Thuwal, Saudi Arabia

Peng Zhan
peng.zhan@kaust.edu.sa
King Abdullah University of Science and Technology
Thuwal, Saudi Arabia

ABSTRACT

The creation of circular economies calls for automated processes of handling goods and natural resources, including waste. Accurate knowledge about disposed waste types facilitates their smart processing and maximises the reuse potential. This requires components such as precise sensing, an appropriate architecture to gather sensed data and gain insights with low latency, and brokering tools to connect waste collectors with processors and collection points. This paper describes a smart waste detection prototype that has been designed and assembled to investigate the automated processes with those components. Its design is centered around a continuum computing architecture to align energy-efficient edge sensing with the shared aim of circular economies to reduce the ecological footprint.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**; • **Computer systems organization** → *Cloud computing*; • **Computing methodologies** → *Machine learning*.

KEYWORDS

cyber-physical applications, circular economy, object detection

ACM Reference Format:

Josef Spillner, Mengwei Liu, and Peng Zhan. 2021. Demo: Smart Waste Disposal with Edge-Cloud Continuum Architecture. In *11th International Conference on the Internet of Things (IoT '21), November 8–12, 2021, St.Gallen, Switzerland.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3494322.3494357

1 INTRODUCTION

Smart waste is a term encompassing technology-supported improvement of waste handling, implying a higher degree of awareness, automation and digitalisation. Examples include smart bins that sense and transmit their fill level and detect anomalies such as fires in the vicinity [1], waste collection route optimisation [3], and waste type detection through pictures taken by a smartphone [5]. In essence, smart waste systems are cyber-physical systems that are combining hardware with software, often also with cloudware to become practical at larger scale. Smart waste handling is an important contribution to the establishment of circular economies, preserving precious natural resources and lowering the ecological footprint of human civilisation. The fallacy is in offsetting the improved footprint by non-smart heavy-weight computation. Instead, green and sustainable computing models need to be applied to the digital transformation of waste handling systems.

Recent progress in greener cloud computing models has already led to visions and proposals such as sustainable serverless computing that favours certain combinations of edge and cloud computing, in particular for smart or intelligent computations that require machine inference of knowledge based on sensed data [7]. Serverless computing is often realised through a function-based architecture, with functions running at the edge, in the cloud or within message brokers. Emphasising the location of code, *Continuum architectures* provide a higher abstraction level, hiding the physical location of where code (e.g. functions) is running and instead placing code wherever appropriate, close to data and constrained by resource capacities and non-functional application requirements such as energy efficiency and latency. Hence, continuum architectures underpinning smart waste handling systems should be investigated and evaluated in practice.

In this paper, we propose an archetypical continuum architecture for the application domain of smart solid waste treatment, in particular the disposal step that initiates the treatment. The architecture is demonstrated as part of a cyber-physical system assembly consisting of several sensors in addition to edge devices and cloud resources. We argue that deploying such systems at larger scale and networking them is feasible and will contribute to a digitalised circular economy.

2 APPLICATION SCENARIOS

Solid waste handling flows depend on the type of product and material. For instance, PET bottles might be collected in central collection points for shredding into flakes and subsequent recycling. PET packaging for food not in bottle form, as well as non-PET bottles, traverse other collection and processing flows. Through sensing the product packaging and material types in conjunction with rules and context information, such as which recycler would buy which quantities of which materials, flows can be automated. Fig. 1 shows exemplary disposal flows and outlines at which points the flows could be augmented with smart waste disposal based on sensing, type detection and subsequent software-defined steps.

The proposed architecture can be used in three distinct scenarios in early stages of the flows for distinguishable solid waste, either at

IoT '21, November 8–12, 2021, St.Gallen, Switzerland

ACM ISBN 978-1-4503-8566-4/21/11.
https://doi.org/10.1145/3494322.3494357

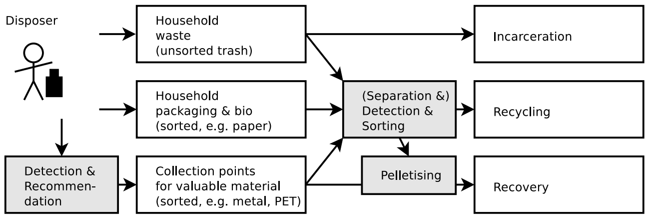

Figure 1: Smartly augmented solid waste flows

disposal time or in processing plants that typically accomplish the requirement through mechanical object separation means ahead of the object detection. We assign the following five functionalities to the scenarios that should be supported by the architecture: *Training* a smart waste packaging and materials database, *inference* of waste classifications from fused sensor data, *reporting* of disposal activities and status, *recommendation* to disposers for whom the rules are often too complex to understand, and *brokering* of post-processing activities (recycling, recovery, incarceration) to bring the waste back into the economical cycle.

The three distinct scenarios covering the functionalities are:

Interactive smart bin. Smart bins take individual items of solid waste. Their interactive augmentation would allow disposers to scan items and receive recommendations, for instance on whether the bin is correct for the selected waste item or whether the bin is already full and another bin should be used instead. Interactive bins combine the *inference* and *recommendation* functionalities ahead of the actual disposal step, and can optionally support non-interactive *training* and *reporting* as well. Moreover, *brokering* within sustainable and circular economies fits in to advise the disposing person of alternative options, such as upcycling for an object detected as valuable.

Non-interactive smart bin. In these bins, waste items are detected as they are disposed. Hence, incorrect disposal is not prevented but could still be notified, for instance acoustically or visually. In non-intrusive mode, they would however not interact with the disposer and instead would provide information only to the operator. Non-interactive bins use embedded sensing to accomplish the *training* and *reporting* functionalities.

Sorting centre. Sorting and recycling centres use conveyor belts that require fast detection of waste through the *inference* functionality, but can also be used for controlled *training*. Moreover, sorting centres critically depend on *brokering* to permit further waste treatment flows.

Owing to the emerging use of robotic sorting facilities, recent research on cloud robotics suggests that the continuum can be extended to robotic arms that autonomously detect objects and their poses and subsequently grab and sort them [2].

3 SYSTEM DESIGN AND ARCHITECTURE

A continuum computing architecture is able to shift computation between edges and clouds as needed, depending on requirements such as resource capacities, resource allocation cost, energy efficiency and non-functional application requirements, primarily related to deadlines and results quality. For the smart waste scenarios, this shifting applies primarily to the compute-intensive machine learning tasks of training and inference, but also to the fusion of sensor data.

To optimise the non-functional properties, we divide the edge into a permanently running p-Edge and an on-demand activated o-Edge, all capable of executing containers that are portable and can be placed and activated where needed. Fig. 2 summarises the architecture in the interactive smart bin flavour. Apart from the involvement of screen and motion sensor, the architecture is reusable for the other scenarios, including for high-throughput detection of objects on a conveyor belt.

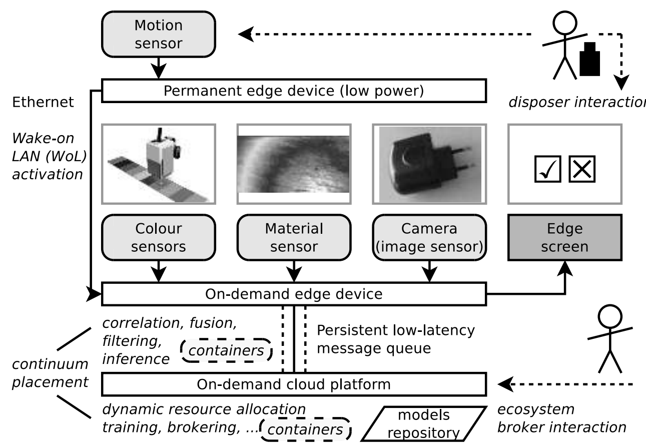

Figure 2: Smart waste treatment continuum architecture

The two key research challenges are:

(1) How can the accuracy of object detection be maximised? This challenge suggests a weighted fusion of multiple sensors, taking their internal accuracies and false positive rates into account. There are practical hurdles, including the size of some sensors such as thermographic surface and coating detection, but also advantages, such as overcoming the insufficient depth of 2D images with physical characteristics sensing [4]. Our prototype allows for investigating the issues concerning real-time fusion, correlation of objects in data streams, weighting and other known problems.
(2) How can the processing time and energy consumption be minimised? Researchers have recently proposed adaptive placement of computation across continuums that take resource capacities into account and learn from past executions or follow stochastic models [6] but practical scenarios to validate them are still needed. Our prototype fills this gap for the narrow domain of object sensing and the wider domain of smart waste treatment.

Further challenges emerge from the engineering of such systems and the incentivisation and integration into sustainable and circular economies, for instance by appropriate bidder interfaces that might require governance for large-scale deployments.

4 SYSTEM IMPLEMENTATION

4.1 Prototype

The prototype focuses on the *interactive smart bin* scenario and hence implements all four functionalities of *training*, *inference*, *reporting* and *recommendation*. Additionally, it implements *brokering* at a small scale.

Hardware. Motion is detected with the Phidget motion sensor MOT2002 connected to a VINT hub that provides network access. For the object detection, three sensor data streams are fused and correlated – Lego EV3 colour sensor, Pixy2 hue tracking colour sensor, and a videocamera. Due to the size of material sensors working on thermal pulses, that sensor is not included in the setup itself although the heatmaps and spectrums generated by it can also be processed within the same system. The p-Edge is implemented with a Raspberry Pi 4 device, whereas a Wake-on-LAN (WoL)-capable notebook serves as o-Edge as well as edge screen to allow for interactive demonstrations. An OpenStack cluster serves as cloud platform, offering the programmatic on-demand allocation of VMs that in turn facilitate containerised function execution and data persistence for both the brokering and the model training.

Software. The edge-hosted software for the interactive bin consists of an on-screen web frontend, a Python backend application with embedded web server and an SQlite-based transactional message queue. Moreover, software agents fill the message queue with discrete sensor measurements or, in the case of the video camera, with information about detected objects via ImageAI and MNetv2. These agents are implemented in the form of event-driven functions. Another software agent manages the connection information and credentials to OpenStack and provides elasticity. Fig. 3 conveys the web frontend view, including the economic effects of waste processors bidding for quantities of materials. To augment the quality of sensor information especially in embedded mode, artificial flash lighting is generated through either the notebook screen or power-controlled USB lights. The p-Edge device runs Etherwake to wake up the o-Edge, a functionality that could in principle also be integrated into the VINT hub itself to reduce the hardware base.

Figure 3: Information screen for disposal system operators

Assembled system. The assembled prototype for a smart waste bin connects the hardware components to an actual multi-material recycling station consisting of several bins, while providing an interactive set-top functionality as well as embedded sensors. Fig. 4 gives a visual impression of the recycling station (p-/o-Edge).

Figure 4: Recycling station setup of the 'interactive smart bin' scenario

The system can also be used in parts for education, with focus on both the circular economy aspects and the computer science aspects. Fig. 5 shows a simplified 'smart bin' scenario implemented on the basis of an EV3 robotic platform with continuum-connected colour sensor and Pixy2 camera for use in classrooms.

Figure 5: Exemplary teaching setup of the 'interactive smart bin' scenario with only the two colour sensors in operation

4.2 Evaluation

Current hardware and software technology is sufficient for adaptively and energy-efficiently placing the machine learning functions across the continuum and for realising the smart waste bin scenario. This finding results from the system evaluation conducted on the continuum spanning the p-Edge (ARMv8 Cortex-A72 CPU @ 1.5

Table 1: Time metrics (in s) related to determination of waste types at edge and cloud, based on 200 images MNetv2 model

Process step	p-Edge time	o-Edge time	Cloud time
Motion sensing	–	0.10	-"-
Edge activation	–	0.95	-"-
Waste sensing	1.14	-"-	-"-
Transmission	–	0.17	0.36
Inference	7.12	4.05	4.97
Total	8.26	6.41 (best)	7.52

GHz), o-Edge (4x i7-5600 CPU @ 2.6 GHz) and cloud (8x Broadwell vCPU @ 2.5 GHz), without GPU acceleration enabled. The dominant contributor to the compute time is the inference based on two image classes with 100 trained images each, including a 90/10 split for five rounds of training and subsequent testing, based on the ImageAI image classification and Tensorflow 2.4. The averaged durations are broken down in Table 1. In contrast, the metrics also suggest that a real-time implementation in hardware for the vision and inference part is required for high-volume detection, for instance in sorting centres.

An additional experiment, conducted synthetically on a lab notebook (i7-8850H CPU @ 2.6 GHz), supports the hypothesis that a smart distribution between continuum resources needs to be conducted, in particular due to the highly diverging compute intensities of training and inference. As indicated by Fig. 6, the duration of the inference steps do not depend on the complexity of the underlying knowledge base derived from training, whereas the training itself is highly depending on both the number of image classes (i.e. types of waste to be recognised) and the number of images per class (i.e. accuracy of visual waste detection). Thus, in a practical setting, a self-trained system for instance based on inputs in the interactive or non-interactive smart bins would have to consider centralised or federated learning instead of learning directly at the edge in order to maintain responsiveness.

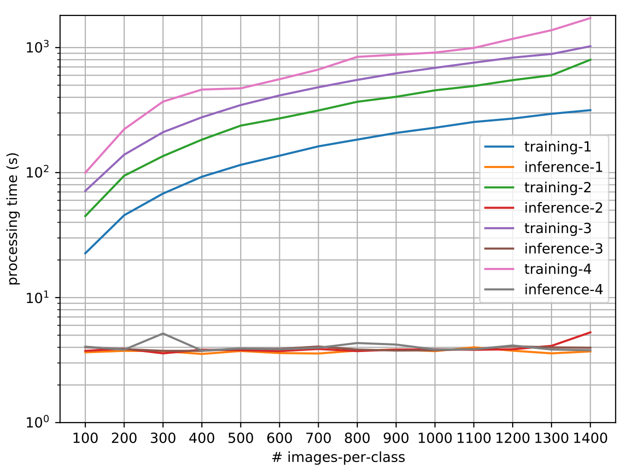

Figure 6: Influence of dataset complexity (1–4 image classes) on training and inference

5 APPLICATION ADOPTION PROSPECTS

Assuming the successful tackling of the two mentioned computer science challenges on generalised sensor data fusion and continuum placement, questions remain on the applicability of the results on industrial scale. Recent market developments suggest that automation and digitalisation levels are increasing throughout the industry. Due to the emergence of specialised hardware in cloud and edge systems, including GPUs and programmable FPGAs, a conceptual refinement of continuum computing concepts to cover transparent acceleration of functionality is highly desired, and even crucial for high-volume waste processing that also leads to higher data volumes. Solid waste sorting systems such as the startup Nommas[1] address the necessary input data quality problem. They combine multiple sensors and vision approaches (HDR camera, ToF lidar, hyperspectral camera and raytracing-supported lens [8]) with deep neural networks and a fleet of robotic delta arms to identify, pick and sort waste objects of relevance. The discretisation and fusion of sensor information leads to the requirement of real-time continuums, including real-time scheduling of edge and cloud services. While large-scale waste processing plants are already available in production and achieve 90–95% purity with purely mechanical means, processing 140kt/year[2], the proliferation of software-driven continuums will allow for lowering cost while increasing purity. More countries adopt the notion of automated sorting of household waste, including Switzerland[3], increasing the target market for technology transfer. The long-term goal of achieving nation-scale circular economies also in countries with currently less developed recycling systems becomes thus viable.

REFERENCES

[1] David Baldo, Alessandro Mecocci, Stefano Parrino, Giacomo Peruzzi, and Alessandro Pozzebon. 2021. A Multi-Layer LoRaWAN Infrastructure for Smart Waste Management. *Sensors* 21, 8 (2021). https://doi.org/10.3390/s21082600
[2] Nicolas Borla, Fabian Kuster, Jonas Langenegger, Juan Ribera, Marcel Honegger, and Giovanni Toffetti. 2021. Teat Pose Estimation via RGBD Segmentation for Automated Milking. *CoRR* abs/2105.09843 (2021). arXiv:2105.09843 https://arxiv.org/abs/2105.09843
[3] Richard Burton Watson and Peter John Ryan. 2021. Visualization and Waste Collection Route Heuristics of Smart Bins Data Using Python Big Data Analytics. In *2021 The 4th International Conference on Software Engineering and Information Management* (Yokohama, Japan) *(ICSIM 2021)*. Association for Computing Machinery, New York, NY, USA, 124–130. https://doi.org/10.1145/3451471.3451492
[4] Kuan-Hung Chen, Yin-Tsung Hwang, and Chih-Peng Fan. 2021. Sensor Data Fusion of Intelligent Autonomous Mover for Object Detection and Collision Avoidance in Environments with Surrounding Crowds. In *2021 IEEE International Conference on Consumer Electronics (ICCE)*. 1–4. https://doi.org/10.1109/ICCE50685.2021.9427728
[5] Alireza Hassani, Alexey Medvedev, Arkady Zaslavsky, and Annamalai Saravanan. 2020. IoT-enabled Smart Waste Disposal System: a Use Case for the Context Management Platform. In *2020 Global Internet of Things Summit (GIoTS)*. 1–6. https://doi.org/10.1109/GIOTS49054.2020.9119592
[6] Petar Kochovski, Rizos Sakellariou, Marko Bajec, Pavel Drobintsev, and Vlado Stankovski. 2019. An Architecture and Stochastic Method for Database Container Placement in the Edge-Fog-Cloud Continuum. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 396–405. https://doi.org/10.1109/IPDPS.2019.00050
[7] Panos Patros, Josef Spillner, Alessandro Papadopoulos, Blesson Varghese, Omer Rana, and Schahram Dustdar. 2021. Towards Sustainable Serverless Computing. *IEEE Internet Computing* (2021).
[8] Qilin Sun, Congli Wang, Qiang Fu, Xiong Dun, and Wolfgang Heidrich. 2021. End-to-End Complex Lens Design with Differentiate Ray Tracing. *ACM Trans. Graph.* 40, 4, Article 71 (July 2021), 13 pages. https://doi.org/10.1145/3450626.3459674

[1] Robotic recycling: https://nommas.com/recycling-sorting/
[2] Sorting at scale in Berlin: https://www.recycling-funktioniert.de/sortieranlage/
[3] Planned Swiss sorting centre: https://www.swissplasticrecycling.ch/