Chapter 5
Robotic Building as Integration of Design-to-Robotic-Production and -Operation

Henriette Bier ⓘ**, Alexander Liu Cheng, Sina Mostafavi, Ana Anton and Serban Bodea**

Abstract *Robotic Building* implies both physically built robotic environments and robotically supported building processes. Physically built robotic environments consist of reconfigurable, adaptive systems incorporating sensor-actuator mechanisms that enable buildings to interact with their users and surroundings in real-time. These robotic environments require *Design-to-Production and -Operation* (D2P&O) chains that may be (partially or completely) robotically driven. This chapter describes previous work aiming to integrate D2RP&O processes by linking performance-driven design with robotic production and user-driven building operation.

5.1 Introduction

While architecture and architectural production are increasingly incorporating aspects of non-human agency employing data, information, and knowledge contained within the (worldwide) network connecting electronic devices, the question is not *whether* but *how* robotic systems can be incorporated into building processes and buildings (Oosterhuis and Bier 2013). This chapter aims to answer this question by reflecting on the achievements of the *Robotic Building* (RB) team at *Technical University Delft* (TU Delft) and by identifying future steps. The focus is on an architecture that is robotically enabled to interact with its users and surroundings in real-time and the corresponding *Design-to-Production and -Operation* (D2P&O) processes that are (in part or as whole) robotically driven. Such modes of production and operation involve agency of both humans and non-humans. Thus agency is not located in one or another but in the heterogeneous associations between them (Latour 2009).

H. Bier (✉) · A. Liu Cheng · S. Mostafavi · A. Anton · S. Bodea
Faculty of Architecture and the Built Environment, TU Delft, Delft, The Netherlands
e-mail: h.h.bier@tudelft.nl

H. Bier · S. Mostafavi
Dessau Institute of Architecture, HS Anhalt, Dessau, Germany

97

This chapter describes attempts to integrate *Design-to-Robotic-Production* (D2RP) with *Design-to-Robotic-Operation* (D2RO) processes by linking design and production with smart operation of the built environment and by advancing applications in performance optimization, robotic manufacturing, and user-driven building operation.

5.2 Robotic Building

RB relies on interactions between human and non-human or cyber-physical agents not only at design and production level but also at building operation level, wherein users and environmental conditions contribute to the emergence of various architectural configurations. Such physically built robotic environments incorporate sensor-actuator mechanisms that enable buildings to interact with their users and surroundings in real-time (see Fig. 5.1). Their conceptualization and materialization require D2RP&O processes that link design to production and building operation (Fig. 5.2). In this context, design becomes process-instead of object-oriented and use of space becomes time-instead of program-or function-based. This implies that architects increasingly design processes, while users operate multiple time-based architectural configurations (Bier and Knight 2014) emerging from the same physical space that may physically or sensorially reconfigure in accordance to environmental and user-specific needs.

In this context, spatial and ambiental reconfiguration optimises use of space by facilitating changing uses of physically built space within reduced timeframes (Liu Cheng and Bier 2016a, b). Furthermore, it reduces energy consumption by employing passive and active climate control and ensures local ambient customisation. Such spatial and ambiental reconfiguration requires virtual modelling and simulation that interface the production and real-time operation of physically built space (Bier and Knight 2014), thus establishing an unprecedented D2RP&O feedback loop, which is the focus of this chapter.

5.2.1 Design-to-Robotic-Production

Industrial robots have been used in a wide range of production processes since the 70s but only more recently academia and creative industry started to explore their potential in architecture. More than 90 institutions and start-ups employ today industrial robots for either developing 1:1 prototypes of bare structures or building components[1] that are integrated in buildings designed and constructed conventionally. In contrast, D2RP aims to introduce strategies for the integral production of buildings

---

[1]The Robotics in Architecture map (accessed from http://www.robotsinarchitecture.org/map-of-creative-robots) shows that more than 90 institutions and start-ups are using robots worldwide.

addressing all structural, environmental, climatic, programmatic, and user-specific, etc. needs. This implies that the complete building process is taken in consideration in order to identify requirements for the robotic production. The goal is to integrate production aspects from the early stages of design.

Several experiments with optimized additive and subtractive production of computationally derived architectural and structural topologies have been implemented at scales ranging from architectural (*macro*) to componential (*meso*) and material (*micro*) scale. By linking performance-based and generative design methods to robotic manufacturing, D2RP processes establish a feedback-loop between design and production of buildings components at full-scale.
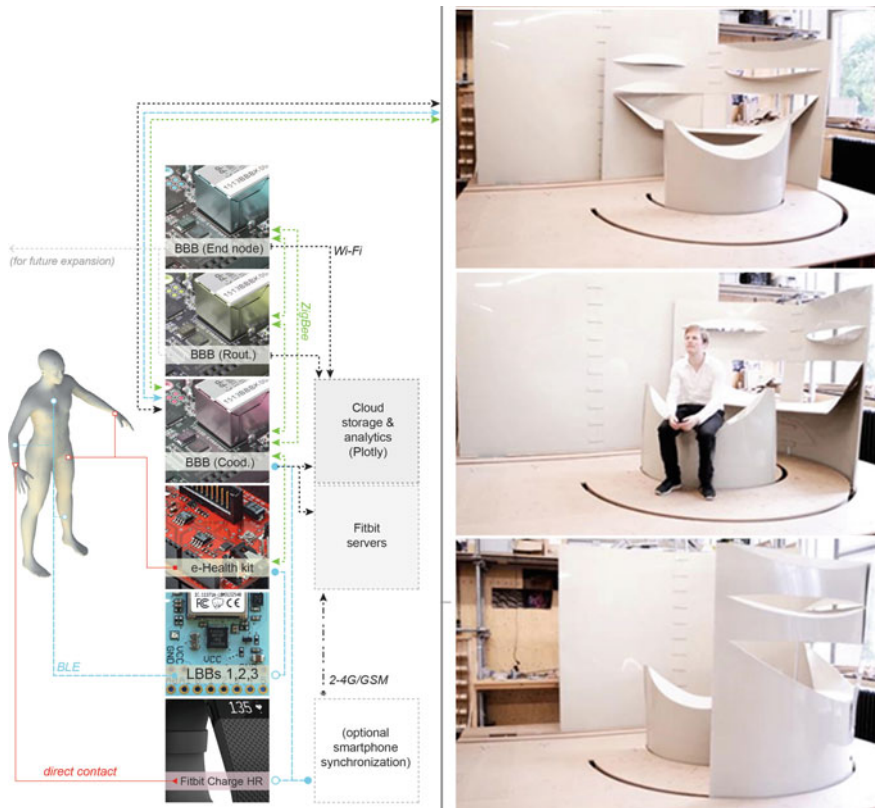

Fig. 5.1 *Design-to-Robotic-Operation* (D2RO) links computational mechanisms and services to spatial reconfiguration for the promotion of occupant well-being (Liu Cheng and Beir 2016a, b). Left: Basic form of the System Architecture, illustrating the relationship between (1) the Local System, (2) the Wearables Subsystem, and (3) the Cloud/Remote Services Subsystem, which are conceived as the essential features of D2RO. Right: *Proof-of-concept* prototype whose physical transformations actuate in response to sensed and processed data

Fig. 5.2 Design-to-Robotic-Production establishing a direct link between virtual modelling and physical fabrication (2014–16). The virtual model (bottom-left) is translated into robotic paths (top) that are further refined using structural analysis in order to robotically produce a clay prototype (bottom-right)

D2RP involves a conversion from the virtual geometric model, which is often the result of optimization processes (e.g. functional, formal, structural, environmental, etc.), into suitable robotic tool paths to deposit, remove, or transform material in order to materialize the intended design. At a digital level, a parametric form-finding approach involving amongst others functional, structural, and environmental optimization is adopted. This approach relies on computational methods such as the *Finite Element Method* (FEM), *Computational Fluid Dynamics* (CFD), etc. Furthermore, material and fabrication constraints are taken into account in order to connect physical materialization with virtual modelling and simulation. This implies that multi-performative design relying on multi-robots production and multi-scale materialisation integrates all requirements from the very beginning of the D2RP process.

5.2.1.1 Multi-performative Computational Design

Architecture is typically developed and built at several discrete scales. While the multi-scalar approach has been the subject of research and debate across architectural history, only more recently—and due to advances in modelling, simulation, and robotic technology—architecture adopted a real-scale design paradigm. In this context, design-to-production focused computation embeds the versatility of computational design into fabrication processes that are accessible to the industry and

designer community. Such integration of computational design into production processes optimizes fabrication resolution, enables novel designs, and promotes a holistic approach in architecture.

Computational design methods developed for D2RP largely rely on recursive computation where once produced geometric results are propagated across design and fabrication iterations resulting in the development of an unified multi-scalar production approach. Consequently, traditional indications of scale—from detail to assembly—and architectural space are translated into wider ranges (*micro*, *meso* and *macro*) that operate as bounds indicative of suitability of particular fabrication techniques and respective recursive depths.

D2RP establishes a feedback-loop between design and fabrication by linking design and simulation environments—e.g., Rhinoceros and Grasshopper—to robotic manufacturing. The role of computation in such robotic production systems is extended, firstly, by the way machines are programmed and, secondly, by the way materials are processed. In the *recursive milling* case study (see Fig. 5.3), continuous robotic paths with embedded information pertaining to material and fabrication constraints generate overall form and surface-texture. The optimised path is a self-avoiding curve[2] that translates into a minimum-length tool-path, featuring low- and high-resolution, for fast and slow material removal.

Since it is particularly suited for delivering designs across multiple scales, recursive milling informs not only subtractive but also additive D2RP. The technology allows access to and control over the internal structure of an object, making the interior of the geometry to design an important subject of research. Variable porosity embodying quantifiable relations between matter and void are employed within D2RP in order to improve environmental performance of building components and reduce material usage. In this context, robotic path-constraints are employed as design drivers to create informed tectonics at volumetric and surface texture levels (see Figs. 5.2 and 5.4). The robotic motion defines the boundaries of the digital design-space in relation to the physical solution-space informing the parametric setup with


Fig. 5.3 Recursive milling method with homogenous resolution (left), tool path with informed resolution based on material removal (middle), and prototyping (right)

---

[2]For example, a Hilbert space-filling curve, which was first described by mathematician David Hilbert in 1891.
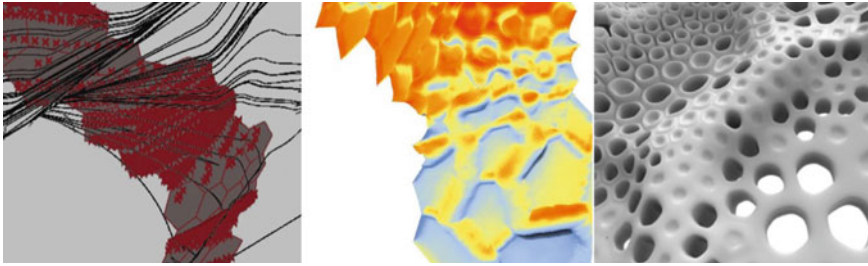
Fig. 5.4 Materialization of informed porosity using structural (left) and environmental analysis (middle) is computed for volumetric tectonics and surface textures

respect to ranges of reachability and optimum tool orientations, thus contributing to enlarging the solution-space.

Structural optimization for additive D2RP, involves methods for form finding of compression-only structures, derived from the innate characteristics of high viscosity ceramic clays. While the optimization takes local and global load and support conditions into consideration, at the *macro* level, a compression-only structure is developed, whose porosity at this scale fulfils functional and aesthetic requirements (Fig. 5.2). At the *micro* level, in order to achieve material porosity, a finite element method for optimizing material distribution is used on selected fragments of the structure (Fig. 5.4). Various algorithmic form finding and optimization techniques, mostly in Rhinoceros–Grasshopper and Python are applied in order to enable the systematic exploration and evaluation of design alternatives within the design-solution space, eventually providing the required information for production.

5.2.1.2 Multi-mode and -Robot Production

As part of a larger D2RP&O framework, D2RP is aiming at integrating the design, fabrication, and operation of buildings in order to address the increasing interest of the construction industry in automation at both production and building operation level. At its core the D2RP system has a cyber-physical setup wherein fabrication sequences are informed by design iterations and simulated kinematic processes. This integrated D2RP approach meeting various demands of the built environment is oriented towards informing building construction processes.

In the workshop at *InDeSem 2015* organised at TU Delft, a compact multi-mode and -robot production setup—consisting of three industrial robots equipped with various tools—was installed in one day to address a large array of manufacturing tasks. Most importantly, these industrial robots were linked directly to computational design environments (Mostafavi and Bier 2016). Once this connection was established, even users at beginner level—e.g., students that have never programmed an industrial robot before—were able to effectively asses fabricability of their designs and optimize the iterations for milling and hot-wire cutting of *Expanded Polystyrene*

(EPS) foam (Mostafavi et al. 2015). Volumetric cutting was used for material removal and general shaping of components while milling was used for adding surface texture and controlling porosity. Such a multi-robot production setup—a de facto small production line—ensures increased efficiency of production while relying on interaction between human and robot agents.

In addition to subtractive D2RP, additive methods were explored where the reach and reduced weight of industrial robots in the small-medium range makes the easily adjustable production unit perfect for the production of small-medium building components. Furthermore, self-developed end-effectors were used for best results in the controlled deposition of customized materials according to patterns that resulted from the structural and robotic path optimization routines (see Fig. 5.5). In this context, the innovation lies in printing with customized materials and end effectors on customized substrates. The robotic setup is flexible enough to allow for the programming of custom paths so that previously fabricated EPS substrates can be used to produce flat or curved ceramic clay pieces (Pottmann et al. 2012). This 3D printing technique is reaching *Technology Readiness Level* (TRL) 6 and could be tested now in an operational environment.

5.2.1.3 Multi-scale Materialization

D2RP employs various materials and relies on multiple robotic production methods in order to achieve quantifiable design performances. Until now materiality, as interface between digital design space and physical fabrication, has been mainly defined along three performance criteria: spatial functionality, structural strength, and environmental efficiency. Furthermore, by integrating computation and robotic materialization, D2RP introduces strategies for extending the design space. With computation implemented at multiple scales and with multi-robot setups enhanced by multi-mode techniques the design space is enlarged. Such multi-mode techniques


Fig. 5.5 Customized 3D printed end-effector and materials for robotic additive manufacturing using ceramic clay

Fig. 5.6 Customized robotic setup for 3D printing of optimized ceramic clay patterns on ruled (left) and flat (right) surfaces

may involve hybrid production approaches that integrate multiple methods of processing materials as for instance, subtractive, additive, and formative.

Considering the building scale, assembly methods allowing expansion beyond the size of building components, which are limited by the actual size of the production space, enlarge the design space as well. These may involve material handling i.e. feeding the components to the robot, picking/gripping and assembling/joining components, while using force control and control of chained tolerances, etc. If multi-robots and -modes processes have been explored since 2014, assembly methods still need to be developed. Thus, multi-scale materialization scenarios, wherein different manufacturing and assembly operations are combined need to be now explored and advanced in order to push the bare structural prototype towards becoming a building.

If the until now developed multi-performative, -mode, and -robot D2RP reaches TRL levels ranging between 4 and 6, multi-scale D2RP is still at its very beginning (Fig. 5.6). Robotisation in building construction by translating building and material performances from discretized geometry into continuous optimized robotic paths (for material deposition, subtraction, or transformation) and by developing coordination scenarios for multi-robot operations in order to involve several robots in the process of production either simultaneously or in short sequence still requires research. In particular the integration between D2RP and D2RO requires further definition, since D2RP pursues robotisation in building construction, while D2RO aims to achieve robotisation in the operation of buildings.

5.2.2 Design-to-Robotic-Operation

Discussions on *intelligence integrated into the built-environment* began in the late 60 s and early 70s (Cook 1970, 1972; Eastman 1972; Pask 1975a, b; Negroponte 1947, 1975). They belonged to a broader discourse that engaged various domains and disciplines in the exploration of opportunities entailed by the *Information Age*. During this period, and partly due to the novelty of the exploration as well as to

the rudimentary state and forbidding costs of *Information and Communication Technologies* (ICTs), these discussions were principally theoretical and/or hypothetical in nature. Over the next two decades, the discourse specialized into subset fields broadly coalescing into the technical on the one hand and the architectural on the other.

With respect to the technical, *Ambient Intelligence* (AmI) was coined in the late 90 s to describe a cohesive vision of a future *digital living room*, a built-environment whose computing hardware and software technology imbued its dwelling space with serviceable intelligence to the benefit of its occupant(s) (Zelkha et al. 1998). Within AmI a further specialized domain developed, i.e., that of *Ambient Assisted Living*—or *Active and Assisted Living*, as preferred by the European Union—(AAL), which framed its inquiry around the promotion of quality of life as well as the prolongation of independence with respect to *Activities of Daily Living* (ADLs) among the elderly via technical assistance. By the first decade of the 21st century, AmI and AAL were established and proliferating topics within the fields of Computer Science and related Engineerings (Augusto et al. 2010; Esch 2013; Cook et al. 2009; Nakashima and Aghajan 2010), Architectural Engineering (Bock et al. 2015; Georgoulas et al. 2014; Linner et al. 2012), and—indirectly—in the Medical Sciences (Acampora et al. 2013).

With respect to the architectural, and beginning with Price's pioneering *Generator Project* and corresponding programs by Frazer and Frazer (1979) in the late 70s, notions of interaction between human and non-human agents in the built-environment began to be explored. For example, in Price's project, architecture was conceived as a set of interchangeable subsystems integrated into a unifying computer system, which enabled a reconfigurability sensitive to function. More importantly, both Price and Frazer intended for the system itself to suggest its own reconfigurations,[3] denoting non-human agency in the built-environment. Although the *Generator Project* was never realized, it became the de facto first instance of a subset field in Architecture concerned with bi-directional communication and interaction between human and non-human agents in the built-environment, viz., *Interactive Architecture* (IA) (Fox and Kemp 2009; Fox 2010; Oosterhuis 2012) first and *Adaptive Architecture* (AA) (Jaskiewicz 2013; Kolarevic 2014; Schnädelbach 2010) later, which—like AmI—have also proliferated in the 21st century.

The embedding of intelligence into the built-environment with respect to AmI/AAL and to IA/AA has differed in sophistication, with the former far surpassing the latter in terms of technical complexity, reliance, and performance. This has been largely due to their differing emphases, with the technical focusing on computing hardware and software technology and the architectural on spatial experience, materiality, and form. That is, the technical proliferated with resources resulting

---

[3]Steenson quotes (2014) two interesting excerpts from letters exchanged by Price and the Frazers. First, from Price to the Frazers, stating his objective: "*The whole intention of the project is to create an architecture sufficiently responsive to the making of a change of mind constructively pleasurable*" (Price et al. 1978). Second, from the Frazers to Price, expressing a desired characteristic: "*If you kick a system, the very least that you would expect it to do is kick you back*" (Frazer and Frazer 1979).

from robust and sustained computational development over decades in ways that the architectural could not, at least not with the same affinity and immediacy. Nevertheless, technical sophistication or lack thereof alone has not necessarily guaranteed or disqualified contributions in the discourse. Indeed, principally technical as well as principally architectural explorations have both independently identified key *effective* as well as *affective* desiderata common to built-environments—intelligent or otherwise—construed as successful with respect to function as well as to spatial experience. This consideration, however, includes a caveat: while both the technical as well as the architectural have yielded independent contributions, these have been otherwise limited by the lack of mutually provided input and/or feedback. For example, AmI/AAL may continue to proliferate as a technical subject even if the physical aspect of its built context remains presupposed and/or static to conventional design and construction frameworks. Similarly, IA/AA may also continue to proliferate in its affective and/or qualitative explorations even if the technical aspects of its implementations express modest computational sophistication. However, the promise of solutions yielded by both principally technical AmI/AAL and principally architectural IA/AA explorations will be unwittingly and invariably limited by the rigid and increasingly outdated character of their complementing frameworks. This is because the sophistication of a system will depend on that of its mutually complementing subsystems; and two or more subsystems may not mutually complement, sustain, and/or support one another properly if their levels of development and sophistication do not correspond (Milgrom 1990). More succinctly expressed: at present, the architectural does not correspond to the technically superior AmI/AAL, while the technical does not correspond to the architecturally superior IA/AA. Consequently, a different design paradigm/framework is required in order to enable comprehensively and cohesively intelligent built-environments with corresponding levels of technical and architectural sophistication.

In this section, principles and strategies developed at TU Delft are introduced as *Design-to-Robotic-Operation* (D2RO), which is presented and promoted as part of an alternative design and development paradigm (i.e., D2RP&O) of intelligent built-environments that considers the technical as well as the architectural in conjunction from the early stages of the design and development processes. In this manner, the built-environment is construed as a highly sophisticated and integrated *Cyber-Physical System* (CPS) (Rajkumar et al. 2010) consisting of mutually informing computational and physical mechanisms that operate cooperatively and continuously via a highly heterogeneous, partially meshed, and self-healing *Wireless Sensor and Actuator Network* (WSAN) (Yang 2014). Via a series of limited and progressively complex proof-of-concept implementations (Liu Cheng 2016, Liu Cheng and Bier 2016a, b; Liu Cheng et al. 2017; Liu Cheng et al. 2017), the feasibility and promise of D2RO are demonstrated and validated.

The current state and development of D2RO is described in the following seven subsections, the first corresponding to the underlying and enabling ICT framework, and the remaining six to mechanisms and/or features that—in conjunction—service an intelligent built-environment capable of intuitive action, reaction, and interaction as well as proactive intervention. The subsystems detailed in these sections have

been implemented from medium to high TRLs (i.e., 5–9) (European Association of Research and Technology Organisations (EARTO) 2015), and constitute an architecturally limited yet technically integrated whole with an overall system TRL of 5.[4] Accordingly, and while those subsystems with TRL 9 are ready to be deployed within commercial solutions, the overall system continues to be developed further both to higher degrees of TRL as well as to include additional subsystems to expand its capabilities.

5.2.2.1 System Architecture

This system at TRL *9* level consists of the following four subsystems (see Fig. 5.7): (1) a *Local system*, which establishes the WSAN; (2) a set of *Wearables*, which extend network's sensing capabilities to include more personal ranges; (3) *Remote/Cloud Services*, which connect the network with Internet-based services and unctions; and (4) Ad Hoc *Support* interfaces, which enable direct user-interventions within the network.

The main difference of the present architecture from that of existing AmI frameworks/solutions is that its functions are not centered on a locally *structured environment*. Instead, the present system is a subsystem within a larger whole. It is *extended* in terms of both its sensing as well as its actuation capabilities, both of which may perform beyond the local *structured environment*. For example—and with respect to *sensing*—in the present architecture, the *local system* continues to monitor the user's activity levels even when he/she is outside of the local *structured environment*. That is, the user-activity recorded by an activity tracker (see item 9, Fig. 5.7) is downloaded by the *local system* from the tracker's manufacturer's servers via an official *Application Program Interface* (API). This enables the local WSAN to process user-activity data continuously, which is necessary in order to develop high-fidelity personalization (Liu Cheng and Bier 2016a, b). With respect to *actuating*, in a situation where the user has collapsed and is unresponsive, the system is capable of acting beyond its local *structured environment* by sending free as well as fee-based SMS/email notifications to care-takers and/or family-members for intervention purposes (Liu Cheng et al. 2016).

Another difference is that the underlying and enabling WSAN is designed as highly heterogeneous—in terms of hardware, software, and communication protocols—in order to subsume functional, operational, and economic advantages across technologies (see Fig. 5.7). Admittedly, researchers have noted that commercial and/or proprietary solutions are often closed, rendering seamless integration with non-commercial and/or non-proprietary solutions highly cumbersome (at best) or unfeasible (at worst) (Harrison et al. 2010). This has raised challenges related to interoperability within heterogeneous systems (see Jiménez-Fernández et al. (2013),

---

[4]For reasons pertaining to system reliability and robustness, the overall TRL is determined by the least developed subsystem, as the failure of subsystem may compromise the serviceability and performance of the whole.
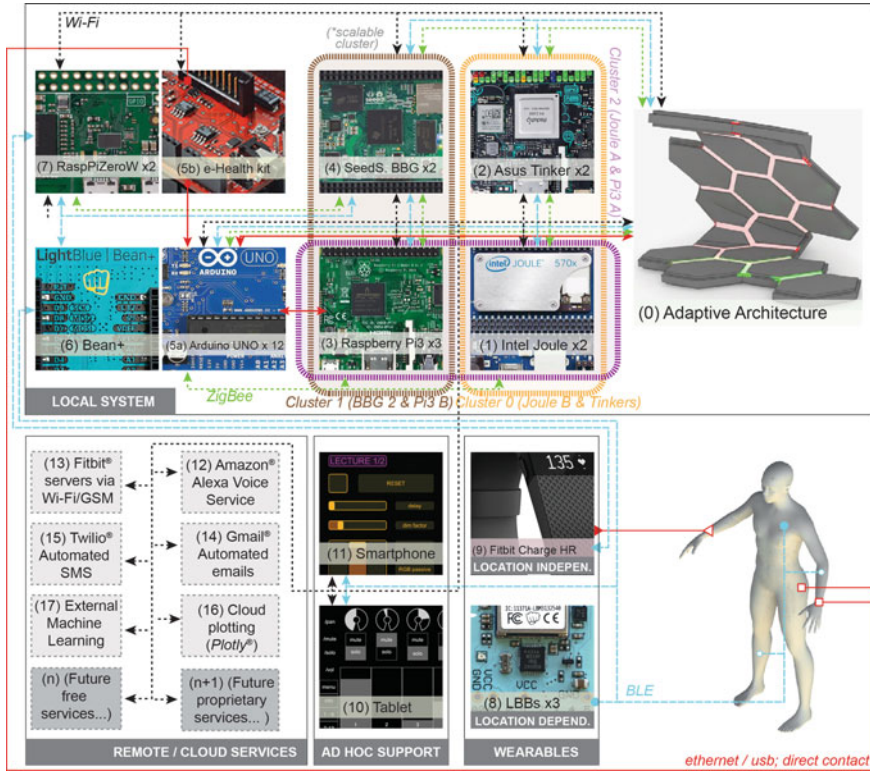
Fig. 5.7 Present state of the *Design-to-Robotic-Operation* (D2RO) System Architecture shown in its basic form, in Fig. 5.1, *Left*. This System Architecture adds a fourth subsystem to the previously identified three subsystems: (1) the Local System; (2) the Wearables Subsystem; (3) the Remote/Cloud Services Subsystem; and (4) the Ad Hoc Support Subsystem

for example), which is partly the reason why some AmI solutions have implemented homogeneous products and/or protocols. Nevertheless, in the last five years manufacturers of proprietary products and services have acted on a vested interest in making their products interoperable with a variety of systems in order to broaden their market. Consequently, an increasing number of proprietary APIs have enabled seamless integration of some proprietary products and services with non-proprietary counterparts.

By virtue of its framework of subsystems as well as of its heterogeneity, the system is highly scalable and open, capable of growing or shrinking to fit a variety of scales and scopes; and of integrating newer devices and of deprecating outdated ones in order to respond more appropriately to evolving tasks at hand.

Subsystem 1: Underlying Mechanism, Local System

This system represents the core of the WSAN. In it a variety of *Microcontroller Units* (MCUs) and development platforms serve as nodes dependent on the local structured environment. Nodes with low-storage and limited information processing capabilities serve as low-energy end devices/routers, and are principally responsible for intermittent sensor-data gathering and relaying.[5] These nodes communicate via BLE in low-range and ZigBee in high-range. Nodes with open storage-capacities, medium-performance information processing capabilities gather and store raw sensor data, parse it, and both make it available to any nodes in the network as well as stream it to Plotly® via WiFi.[6]

Nodes with high-performance information processing capabilities are principally responsible for coordination and computation.[7] These nodes may be clustered to form more powerful nodes depending on the load-requirement and exchange data with one another and with other nodes via WiFi, BLE, or ZigBee depending on the frequency as well as the latency-requirement. In one particular case, wired connections are used between nodes for data exchange (i.e., item 5 with 3, Fig. 5.7). If necessary, all Linux-running devices, regardless of individual computational power or predetermined function, may conform a cluster.

The present configuration is one of possible many. The items featured as well as the multiple instances of each serve to represent a typical highly heterogenous (both in terms of architecture as well as communication protocols and services) and cost-effective foundation capable of sustaining the growing complexity of subsequent developments and implementations.

Subsystem 2: Wearable Devices

A set of three Light Blue Bean™s (LBBs) conform the location dependent wearables while a Fitbit® Charge HR™ activity tracker (item 9, Fig. 5.7) the location independent wearable. The former detects movement in the upper-body, upper- and lower-extremities and advises the system to listen for *Open Sound Control* (OSC) packets corresponding to accelerometer data sent from a smartphone (see subsystem 3 below). Alternatively, if no smartphone is present, the LBBs broadcast accelerometer data via BLE into the system as well. This alternative is relegated to a contingency measure due to the energy-consumption of constant and sustained data streaming. Both OSC and BLE accelerometer data are used to build and update *Support Vector Machine* (SVM) and k-*Nearest Neighbor* (k-NN) classification models and to feed real-time data in the *Machine Learning* (ML) mechanism for *Human Activity Recognition* (HAR).

---

[5] Viz., PunchThrough® Bean+™ and Arduino® UNO™—items 5, 6, and 8, Fig. 5.7.

[6] Viz., Raspberry® Pi Zero W™ (RPiZW)—item 7, Fig. 5.7.

[7] Viz., Intel® Joule™, Asus® Tinkerboard™, Raspberry® Pi 3™ (RPi3) and SeedStudio® BeagleBone Green™—items 1–4, Fig. 5.7.

The principal function of the activity tracker is to gather heart-rate and physical activity (in terms of steps taken and distance covered) data continuously regardless of the location. When the user is inside the structured environment, the LBBs in conjunction with a smartphone also provide user-activity data to the system for HAR. But when the user is outside of the environment, the WSAN continues to draw limited data gathered by the activity tracker by downloading it from Fitbit®'s servers (the tracker synchronizes with the servers via mobile data when WiFi is unavailable).

Subsystem 3: Remote/ Cloud Services

Six cloud-based services conform this subsystem, three of which were first integrated in the ISARC 2016 conference article (Liu Cheng and Bier 2016a, b), and three others newly integrated into the current ecosystem. The inherited three are the following: (I) external ML mechanism via MATLAB® (item 17, Fig. 5.7); (II) data exchange with Fitbit®'s servers via its API (item 13, Fig. 5.7) and (III) cloud data-storage and -plotting via Plotly®'s API (item 16, Fig. 5.7). And the newly integrated three are the following: (*IV*) Amazon®'s AVS (item 12, Fig. 5.7); (V) automated SMS notifications, both via Twilio®'s API (item 15, Fig. 5.7) as well as via a T35 GSM shield as part of one of the end-device nodes of subsystem 1; and (VI) automated email notifications via Gmail©'s API (item 14, Fig. 5.7).

Subsystem 4: Ad Hoc Support Devices

In the last five years, smartphones have become convenient and ubiquitous tools for the tracking of inhabitants across a space (Andò et al. 2014), fall detection (Liu Cheng et al. 2016; Abbate et al. 2012), and HAR via ML (Anguita et al. 2013; Ortiz 2015; Micucci et al. 2017), which in conjunction with their battery life and rechargeability are the principal reasons why it they are the preferred means of accelerometer-data gathering in this development. In addition to this function, a user-interface/configuration mechanism is also enabled via a proprietary (viz., TouchOSC™ by Hexler Limited®) and a free (viz., Control by Charlie Roberts) smartphone application. This mechanism enables the user to override automation by permitting manual input/configuration.

Similarly, a tablet device has also been integrated into the ecosystem in order to provide both another user interface with a more comfortable viewing area as well as a means to modify the behavior of the LBBs and Bean + devices via BLE. Unlike the Linux-based devices of the ecosystem, the LBBs and Bean + cannot be accessed wirelessly via *Secure Shell* (SSH). Nevertheless, any necessary modifications to the devices' program or sketch may be effected wirelessly via the tablet. For example, one of the LBBs could be tasked with gathering temperature data on the user for a certain period of time and at varying intervals instead of notifying acceleration events. Both the smartphone and the tablet may access the LBBs and Bean + devices

via BLE, and both are installed with the user-interface/configuration applications to enable parallel modifications should this be necessary.

5.2.2.2 Global/Local Ventilation Mechanism

This mechanism reaching TRL 5 is first implemented and tested via an abstracted surrogate model equipped with twelve DHT-22 temperature and humidity sensors, twelve air-quality sensors,[8] and twelve small DC-motor fans connected to three RPiZWs and one RPi3.

As corroborated by the *Comité Européen de Normalisation* (CEN) Standard EN15251-2007 (2007) as well as ASHRAE Standard 55-2013 and Standard 62.1-2013 (2013), the thermal Environmental Conditions for Human Occupancy with respect to comfort should be 67 to 82 °F. (~19.5–27.8 °C.) (ASHRAE® Standard 2013), while relative humidity in occupied spaces be less than 65% in order to discourage microbial growth. Furthermore, independent of human comfort considerations, frequent and consistent ventilation reduces the concentration of toxins in the air as well as the prevalence of airborne diseases (2009). In this *proof-of-concept* setup, if the collective temperature or humidity levels exceed recommended limits for comfort, all the fans activate, thereby drawing fresh air into the inhabited space (i.e., Global ventilation concept). If, however, certain areas exceed either or both limits, only those fans within and surrounding them activate (i.e., Local ventilation concept). The same concept holds for instances of air-pollution.

5.2.2.3 Voice-Control Mechanism via *Alexa Voice Service*

This mechanism reaching TRL 9 is implemented and tested via the same RPi3 mentioned in the previous section, an open-source repository using Amazon®'s API (GitHub Inc.© 2017), and a generic microphone as well as repurposed speakers. The flexibility of developing custom—and more affordable—*Alexa-enabled Devices* permits virtually any built-environment device, whether deployed in an architectural or an urban context, to capitalize from AVS.

Two main objectives inform the present integration. The first is to enable a powerful and scalable voice-control mechanism within the present development. The second is to demonstrate a cohesive technological heterogeneity between an open-source WSAN and a proprietary commercial service without additional cost (with respect to Fitbit® and Gmail©) or with minimum cost. This latter consideration connects a local intelligent-built environment with vast resources in the WWW, enabling the user to engage in a variety of activities from streaming music to purchasing groceries via devices fundamentally embedded into the built-environment.

---

[8] Viz., three of each: MQ-3 *Alcohol*, MQ-4 *Methane*, MQ-7 *Carbon Monoxide*, and MQ-8 *Hydrogen Gas*.

In the present state of D2RO, the scope of service of AVS is limited to predefined web-based skills. Work is being undertaken to expand scope to encompass services deployable within the local structured environment by either integrating a growing number of smart-home products compatible with AVS or by creating custom skills to suit specific *Internet of Things* (IoT) open-source devices via ASK.

5.2.2.4 Intervention via SMS and Email Notifications Mechanism[9]

This mechanism at TRL 9 level is implemented and tested via another RPiZW node, a smartphone, and Twilio®'s as well as Gmail©'s APIs. Additionally, a non-web-based contingency device is developed using a Siemens® T35 GSM shield mounted on an Arduino® UNO™. The main objective with this implementation is to setup the foundations of an increasingly comprehensive intervention framework capable of reacting to emergency events, both with respect to the inhabitants of the built-environment and with this environment per se. The Twilio® implementation represents a cost-effective SMS service, while the T35 GSM setup represents a standard prepaid SMS service. A scenario may be entertained where the built-environment's WiFi service is unavailable for a period of time, yet the integrity of the WSAN's core (i.e., the local system) remains uncompromised as its constituents remain networked via ZigBee and BLE. In such a scenario, an emergency event may be reported via the T35 GSM setup, as it relies on standard cellular communication. Conversely, another scenario may also be entertained, where cellular services are unavailable due to lack of coverage. In this scenario, emergency events may be reported via Twilio®'s SMS service to any location worldwide. Both of these hypothetical scenarios presuppose that the recipient is capable of receiving cellular messages at the time of notification. However, this may not be the case. This kind of situation is the motivation behind email notifications. Although it cannot guarantee message reception, it adds yet another means for it. Unlike both SMS mechanisms, the email notification is free.

5.2.2.5 Machine Learning[10]

With respect to the first functionality, a *Machine Learning* (ML) subsystem is integrated in the proposed system-architecture in order to enable *Human Activity Recognition* (HAR) mechanisms (Liu Cheng et al. 2017). With respect to HAR, ML methods have typically used gyroscopic data collected via portable devices (e.g., smartphones, etc.) (Anguita et al. 2013; Ortiz 2015) or via sensor-fusion (Palumbo et al. 2016). The ML subsystem consists of two classification mechanisms developed based on polynomial programming of SVM and *k*-NN classifiers. These SVM and *k*-NN models are built on a dynamically clustered set of high-performance nodes in the localized WSAN.

---

[9]See (Liu Cheng et al. 2016) for a detailed discussion of this mechanism.

[10]See (Liu Cheng et al. 2017) for a detailed discussion of this mechanism.

Due to their evolving and resilient characters, ML classifiers have been implemented in a variety of applications built on WSANs (Alsheikh et al. 2014). HAR, as one such application, has successfully exploited classifiers in the last five years (see, for example, (Xiao and Lu 2015; Villa et al. 2012; Andreu and Angelov 2013). However, due to the cost-effective and low energy-consumption character typical of WSAN nodes, computational processing with respect to feature extraction has been considerably limited (Salomons et al. 2016). To overcome this limitation, the present implementation is capable of instantiating ad hoc clusters consisting of a variety of high-performance nodes. Furthermore, several clusters may be instantiated simultaneously in order to enable parallel high-performance information processing activities.

Another way to overcome this limitation is to avoid it altogether by outsourcing all high-performance information processing to cloud-based ML services.[11] But there are a number of limitations with this approach. The first, and perhaps the most salient, is the cost incurred by including proprietary services in any proposed intelligent built-environment solution. A second yet no less important limitation may be the impact to the solution's resilience. Should the built-environment lose access to the Internet, it would be incapable of generating classification models.

In the current state of D2RO, integration of both cloud-based as well as localized ML capabilities in order to ascertain robustness and resilience. Whenever possible, ML processes are locally and dynamically executed via ad hoc node-clustering. But should this prove impossible either due to failure or unavailability of proper resources, cloud-based ML services are used. More specifically, two ML mechanisms are integrated into the present system: (1) a localized ad hoc cluster system based on open-source and purpose-written *Python* scripts, and (2) a simulated cloud-based analytics service using MathWorks® MATLAB™. In both mechanisms SVM and *k*-NN classification models are generated.

In the localized mechanism, a script based on *pyOSC* is first written to receive OSC data from any device and application capable of broadcasting in described protocol. While all the WiFi-enabled nodes in the system's WSAN have the capacity to receive this data-streaming, only one of the nodes of the cluster instantiated to generate classification models stores it locally and streams it to a cloud-based data visualization service (i.e., Plotly™). Should the receiving node fail, another high-performance node will replace it automatically. Since the proposed solution uses a smartphone and three LBBs for data redundancy, resolution, and validation, the script in question proceeds to parse and to reduce the noise in the received multi-sensor data in order to generate a robust and unified dataset. At this point the dataset is processed through two ML scripts based on *scikit-learn* (Pedregosa et al. 2011; Buitinck et al. 2013), one for SVM and another for *k*-NN classification models.

---

[11]E.g., Google® CloudPlatform™, Amazon® Machine Learning™, Microsoft® Azure™, etc.

5.2.2.6　Object Recognition via OpenCV[12]

The object-recognition mechanism reaching TRL 9 is implemented with open-source BerryNet® (2017), which is built with a classification model (viz., Inception® ver. 3 (Szegedy et al. 2016) as well as a detection model (viz., TinyYOLO® Redmon and Farhadi 2017). The classification model uses *Convolutional Neural Networks* (CNNs), which are at the forefront of ML research (Szegedy et al. 2016). An advantage of BerryNet® is that it is a fully localized DL gateway implementable on a cluster of RPi3 s. On an individual RPi3, the *inference* process is slow, requiring a delay between object-recognition sessions. This situation is ameliorated by the dynamic clustering feature of the WSAN. Another benefit-*cum*-limitation is that BerryNet®'s classification and detection models are pretrained, which avoids the need to generate models locally.

The object-recognition mechanism in D2RO is intended to be deployed across a variety of cameras in the overall built-environment, and that instances of detection were to be cross-referenced to minimize false positives. In order to implement this setup, each RPi3 node in the WSAN is equipped with a low-cost Raspberry Pi Camera® V2.1, then BerryNet® is installed in every node and the *inference* mechanism tested individually. The next step is to enable the nodes to share their detection results, which could be done via WiFi. Nevertheless, in order to reduce energy-consumption for every object-detection cross-referencing instance, ZigBee is preferred. In order to enable ZigBee on BerryNet®'s *detection_server.py* and *classify_server*.*py* were modified and made compliant with *python-xbee* (2017).

5.3　Design-to-Robotic-Production and -Operation

The integration of D2RO with D2RP, as explored at TU Delft, relies on the notion of hybrid componentiality. This implies that components are cyber-physical and their design is informed by structural, functional, environmental, assembly and operation considerations (Mostafavi and Bier 2016). At the *micro*-scale, the material is conceived as a porous system, where the degree and distribution of porosity i.e. density are informed by functional, structural and environmental requirements, while taking into consideration both passive (structural strength, thermal insulation, etc.) and active (adaptive, reconfigurable, etc.) behaviours. At the *meso*-scale, the component is informed mainly by the assembly logic, while at the *macro* scale, the assembly is informed by architectural considerations.

D2RP&O has been explored in the project *Hybrid Assemblies* (see Fig. 5.8) implemented with students *Dessau Institute of Architecture* (DIA). The project focused on the development of architectural systems composed of heterogeneous components addressing various requirements from functional and formal to structural and climatic. While taking these requirements into account, the project focused on the

---

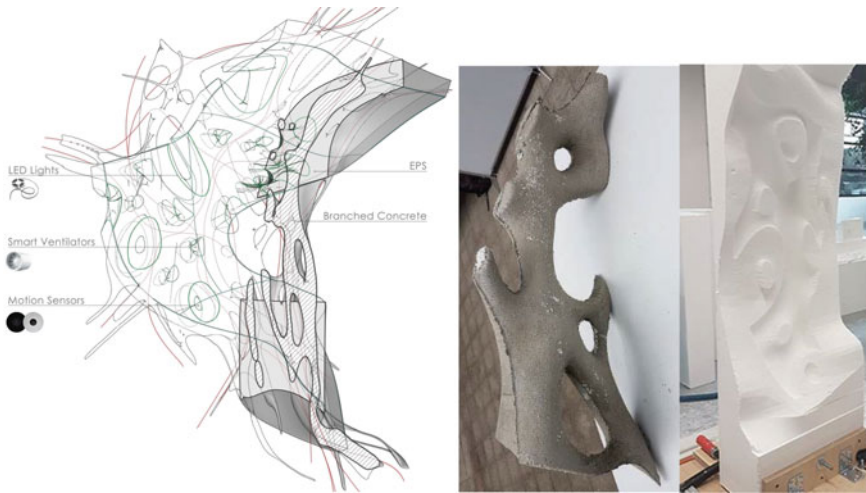[12]See (Liu Cheng et al. 2017) for a detailed discussion of this mechanism.

Fig. 5.8 Multi-layered D2RP&O integration logic (left) of fragment made of concrete (middle) that is cast in robotically produced EPS (right)

notion of embedded interactive or adaptive systems employed for climate control. The distributed, dynamic climate control has been conceived as consisting of intelligent networked climate control components, locally driven by people's preferences and changing environmental conditions. The challenge was to integrate the passive energy saving material architecture with the active climate control that is taking into account changes in the use of space and respective fluctuating needs based (not on average but) on real-time data.

The design was defined by optimization strategies involving spatial configuration, structural analysis, heating and cooling, lighting requirements, and the integration of ICT devices.[13] While, structural analysis is employed to map areas that are needed for structural support, lighting is determined based on 24/7 activities and their corresponding requirements. These inform the shape and the location of cavities for LED-based illumination. Then heating and cooling requirements are identified for the integration of intelligent ventilation systems as well as the required sensors for automated control.

The multi-layered hybrid components consisting of concrete, EPS, and smart devices follow *componentiality* and *hybridity* principles characteristic of D2RP&O. Layers are designed in direct response to a purpose or a function. For example, the concrete layer is formed following the stress lines and cavities in the EPS layer are designed according to ICT-integration requirements (Fig. 5.8). This approach embeds all cyber-physical requirements from the onset of the design process.

With respect to Systems Architecture, the detailed object-recognition mechanism adds another means for the system to become aware of the built-environment. In the

---

[13]See (Liu Cheng et al. 2016) for a detailed discussion of this mechanism.

setup discussed, the deployment scenario is construed as a single-occupant housing unit. But in scenarios with more occupants, the recognition of each individual may instantiate actuations and transformations in the built-environment specific to each individual's preferences.

The integration of D2RP with D2RO as explored at TU Delft is unprecedented in particular because of the focus on buildings. Installations such as Open Columns or the Hyozolic series, which reconfigure according to changing levels of CO_2 or movement of people, may be integrating computational design with additive manufacturing and smart reconfiguration but their application to buildings is still speculative.[14]

5.4   Conclusion

D2RP&O is unique in its aim to link design and production with smart operation of the built environment and advances applications in performance optimization, robotic manufacturing, and user-driven operation in architecture. Relying on human and non-human interaction in the design, production, and operation of buildings, D2RP&O is fundamentally changing the role of the architect. Architects design increasingly processes not objects, while users operate multiple time-based architectural configurations emerging from the same physical space that reconfigures in accordance to environmental and user specific needs. In this context, D2RP&O empowers architects to regain control over the design implementation into physically built environments and allows end-users to participate as co-creators in the adaptation i.e. customization of their environments over time.

Even if D2RP and D2RO have been developed as separate areas of research, their partial integration into a coherent D2RP&O chain has been implemented and tested in the *Hybrid Assemblies* project. This integration indicated that D2RP&O could significantly contribute to improving material-, energy-, and process-efficiency, as well as (structural, environmental, functional, etc.) performance of buildings.

In addition to developing a coherent D2RP&O chain, the challenge for the future is the integration of Human-Robot Interaction (HRI). For instance, by employing laser scanning to capture the current status of building process, an extended feedback-loop between the virtual and the physical environments is established. D2RP robots may then interact with humans, as for instance, human operators may teach robots to do certain tasks by guiding them with a tool or by hand, while dynamic safety systems are in place,[15] etc. Similarly, D2RO relies on HRI when sensor-actuators ensure that inhabitants can customize the use of the physically built space. Main consideration

---

[14]The two installations were developed as architecture inspired art projects (accessed from http://cast.b-ap.net/opencolumns/ and http://www.philipbeesleyarchitect.com/sculptures/0929_Hylozoic_Ground_Venice/).

[15]HRI is in detail described in the chapter titled "Human-Robot Collaboration and Sensor-Based Robots in Industrial Applications and Construction" of this volume.

is that production and operation of buildings will be in the future robotized and identifying which skills sets are better acquired and executed by humans while others by machines is key to developing interaction scenarios between humans and robots.

References

Abbate S, Avvenuti M, Bonatesta F, Cola G, Corsini P, Vecchio A (2012) A smartphone-based fall detection system. Pervasive Mobile Comput 8(6):883–899

Acampora G, Cook DJ, Rashidi P, Vasilakos AV (2013) A survey on ambient intelligence in healthcare. Proc IEEE 101(12):2470–2494

Alsheikh MA, Lin S, Niyato D, Tan H-P (2014) Machine learning in wireless sensor networks: algorithms, strategies, and applications. IEEE Commun Surv Tutor 16(4):1996–2018

Andò B, Baglio S, Lombardo CO, Marletta V (2014) An advanced tracking solution fully based on native sensing features of smartphone. In: 2014 IEEE Sensors Applications Symposium (SAS), pp 141–144

Andreu J, Angelov P (2013) An evolving machine learning method for human activity recognition systems. J Ambient Intell Hum Comput 4(2), 195–206. https://link-springer-com.ezproxy.library.ubc.ca/content/pdf/10.1007%2Fs12652-011-0068-9.pdf

Anguita D, Ghio A, Oneto L, Parra X, Reyes-Ortiz JL (2013) A public domain dataset for human activity recognition using smartphones. In: Proceedings of the 21th european symposium on artificial neural networks, computational intelligence and machine learning, ESANN 2013

ASHRAE® Standard 62.1-2013, 2013

ASHRAE® Standard 55-2013, 2013

Augusto JC, Novais P, Corchado JM, Analide C (eds) Ambient Intelligence and Future Trends: International Symposium on Ambient Intelligence (ISAmI 2010). Advances in Intelligent and Soft Computing, vol 72. Springer Verlag

Bier HH, Knight T (2014) Digitally-driven architecture (English). Footprint 4(1): 1–4

Bier HH, Knight T (2014b) Data-driven design to production and operation. Footprint 8(2):1–8

Bock T, Güttler J, Georgoulas C, Linner T (2015) The development of intra-house mobility, logistics and transfer solutions in PASSAge. JRM 27(1):108

Buitinck L, Louppe G et al (2013) API design for machine learning software: experiences from the scikit-learn project. http://arxiv.org/pdf/1309.0238

Comité Européen de Normalisation© (CEN) (2007) Standard EN 15251–2007: Indoor environmental input parameters for design and assessment of energy performance of buildings addressing indoor air quality, thermal environment, lighting and acoustics http://www.sysecol2.ethz.ch/OptiControl/LiteratureOC/CEN_06_prEN_15251_FinalDraft.pdf. Accessed 16 July 2007

Cook P (1970) Experimental architecture. Universe Books, New York

Cook P (1972) Archigram. Studio Vista, London

Cook DJ, Augusto JC, Jakkula VR (2009) Ambient intelligence: Technologies, applications, and opportunities. Pervasive Mobile Comput 5(4):277–298

DT42©, Ltd, *BerryNet® (2017)* Deep learning gateway on Raspberry Pi. https://github.com/DT42/BerryNet. Accessed 21 Jun 2017

Eastman CM (1972) Adaptive conditional architecture. Institute of Physical Planning, School of Urban and Public Affairs, Carnegie-Mellon University, Pittsburgh

Esch J (2013) A survey on ambient intelligence in healthcare. Proc IEEE 101(12):2467–2469

European Association of Research and Technology Organisations (EARTO) (2015) The TRL Scale as a Research & Innovation Policy TOOL: EARTO Recommendations. http://www.earto.eu/file admin/content/03_Publications/The_TRL_Scale_as_a_R_I_Policy_Tool_-_EARTO_Recomm endations_-_Final.pdf. Accessed 07 Jan 2015

Fox M (2010) Catching up with the Past: A small contribution to a long history of interactive environments

Fox M, Kemp M (2009) Interactive architecture, 1st edn. Princeton Architectural Press, New York

Frazer J, Frazer J (1979) Letter To Cedric Price, Jan 1979, Box 5, Generator document folio, DR1995:0280:65 5/5, CPA

Georgoulas C, Raza A, Güttler J, Linner T, Bock T (2014) Home environment interaction via service robots and the leap motion controller. In: Proceedings of the 31st international symposium on automation and robotics in construction (ISARC 2014)

GitHub Inc.© (2017) alexa-avs-sample-app. https://github.com/alexa/alexa-avs-sample-app. Accessed 01 Apr 2017

Harrison BL, Consolvo S, Choudhury T (2010) Using multi-modal sensing for human activity modeling in the real world. In: Nakashima H, Aghajan HK, Augusto JC (eds) Handbook of ambient intelligence and smart environments. Springer, New York, pp 463–478

Jaskiewicz T (2013) Towards a methodology for complex adaptive interactive architecture

Jiménez-Fernández S, de Toledo P, del Pozo F (2013) Usability and interoperability in wireless sensor networks for patient telemonitoring in chronic disease management. IEEE Trans Bio-med Eng 60(12):3331–3339

Kolarevic B (2014) Outlook: Adaptive Architecture: Low-Tech, High-Tech, or Both?. In: Kretzer M, Hovestadt L (eds) Applied Virtuality Book Series, v.8, ALIVE: Advancements in adaptive architecture. Birkhäuser, Boston, pp 148–157

Latour B (2009) Reassembling the social: an introduction to actor-network-theory. Oxford University Press, Oxford

Linner T, Georgoulas C, Bock T (2012) Advanced building engineering: Deploying mechatronics and robotics in architecture. Gerontechnology 11(2):380

Liu Cheng A (2016) Towards embedding high-resolution intelligence into the built-environment. Archidoct 4(1), 29–40 http://www.enhsa.net/archidoct/Issues/ArchiDoct_vol4_iss1.pdf

Liu Cheng A, Bier HH (2016a) An extended ambient intelligence implementation for enhanced human-space interaction. In: Proceedings of the 33rd international symposium on automation and robotics in construction (ISARC 2016), pp 778–786

Liu Cheng A, Bier HH (2016b) Adaptive building-skin components as context-aware nodes in an extended cyber-physical network. In: Proceedings of the 3rd IEEE world forum on internet of things. IEEE, pp 257–262

Liu Cheng A, Bier HH, Mostafavi S (2017) Deep learning object-recognition in a Design-to-Robotic-Production and -Operation implementation. In: Proceedings of the 2nd IEEE Ecuador technical chapters meeting 2017, Guayaquil, Ecuador

Liu Cheng A, Georgoulas C, Bock T (2016) Fall detection and intervention based on wireless sensor network technologies. Autom Constr

Liu Cheng A, Bier HH, Latorre G, Kemper B, Fischer D (2017) A high-resolution intelligence implementation based on Design-to-Robotic-Production and -Operation strategies. In: Proceedings of the 34th international symposium on automation and robotics in construction (ISARC 2017)

Micucci D, Mobilio M, Napoletano P, Tisato F (2017) Falls as anomalies?: An experimental evaluation using smartphone accelerometer data. J Ambient Intell Humaniz Comput 8(1):87–99

Milgrom PR (1990) The economics of modern manufacturing: technology, strategy, and organization. Am Econ Rev 80(3):511–528

Mostafavi S, Bier H (2016) Materially informed design to robotic production: A Robotic 3D printing system for informed material deposition. In: Reinhardt D, Saunders R, Burry J (eds) Robotic fabrication in architecture, art and design 2016. Springer, Heidelberg, pp 339–349

Mostafavi S, Bier H, Bodea S, Anton AM (2015) Informed design to robotic production systems; developing robotic 3D printing system for informed material deposition. In: Proceedings of the 33rd eCAADe conference real time, vol. 2, Vienna (Austria), 16–18 Sept 2015 https://repository.tudelft.nl/islandora/object/uuid%3A02bf0811-41f0-4ac9-a270-af796fc7 7c7c/datastream/OBJ/download

Nakashima H, Aghajan HK, Augusto JC (eds) (2010) Handbook of ambient intelligence and smart environments. Springer, New York

Negroponte N (1969) Toward a theory of architecture machines. J Archit Educ (1947–1974) 23(2): 9

Negroponte N (1975) The architecture machine. Comput Aided Des 7(3):190–195

n.io Innovation© (2017) LLC, python-xbee: Python tools for working with XBee radios. https://gi thub.com/nioinnovation/python-xbee. Accessed 01 Jun 2017

Oosterhuis K (2012) Hyperbody: First decade of interactive architecture. Jap Sam Books, Heijningen

Oosterhuis K, Bier HH (2013) IA #5: Robotics in architecture. Jap Sam Books, Heijningen

Ortiz JLR (2015) Smartphone-based human activity recognition. Springer, Cham

Palumbo F, Gallicchio C, Pucci R, Micheli A (2016) Human activity recognition using multisensor data fusion based on Reservoir computing. AIS 8(2):87–107

Pask G (1975a) Conversation, cognition and learning: A cybernetic theory and methodology. Elsevier, Amsterdam

Pask G (1975b) The cybernetics of human learning and performance: A guide to theory and research. Hutchinson Educational, London

Pedregosa F, Varoquaux G et al (2011) Scikit-learn: Machine Learning in Python. http://arxiv.org/ pdf/1201.0490

Pottmann H, Asperl A, Hofer M, Kilian A (2012) Architectural geometry. Bentley Institute Press, Exton

Price C (1978) Letter To John Frazer, Dec 1978, Box 4, Generator document folio, DR1995:0280:65 4/5, CPA

Redmon J, Farhadi A (2017) YOLO9000: Better, Faster, Stronger. http://arxiv.org/pdf/1612.08242

Salomons EL, Havinga PJM, van Leeuwen H (2016) Inferring human activity recognition with ambient sound on wireless sensor nodes. Sensors (Basel, Switzerland) 16(10)

Schnädelbach H (2010) Adaptive architecture-a conceptual framework. In: Proceedings of Media City

Seppänen O, Kurnitski J (2009) Moisture control and ventilation. World Health Organization

Steenson MW (2014) Architectures of information: Christopher Alexander, Cedric Price, and Nicholas Negroponte and MIT's Architecture Machine Group

Szegedy C, Vanhoucke V et al (2016) Rethinking the inception architecture for computer vision. http://arxiv.org/pdf/1512.00567

Villa AE (2012) Artificial neural networks and machine learning– ICANN 2012. In: 22nd International conference on artificial neural networks, Lausanne, Switzerland, 11–14 Sept 2012, Proceedings. Springer, New York

Xiao W, Lu Y (2015) Daily human physical activity recognition based on kernel discriminant analysis and extreme learning machine. Math Probl Eng 2015(1):1–8

Rajkumar R, Lee I, Sha L, Stankovic J (eds) Cyber-physical systems: the next computing revolution. ACM

Yang S-H (2014) Wireless sensor networks principles, design and applications. Springer, London

Zelkha E, Epstein B, Birrell S, Dodsworth C (1998) From devices to 'ambient intelligence': the transformation of consumer electronics (Conference Keynote). In: Digital living room conference