2024 IEEE International Conference on Smart Computing (SMARTCOMP) | 979-8-3503-4994-8/24/$31.00 ©2024 IEEE | DOI: 10.1109/SMARTCOMP61445.2024.00025

iFair: Achieving Fairness in the Allocation of Scarce Resources for Senior Health Care

Modeste Mefenya Kenne*, Prasanna Date**, Ronald T. Eguchi†,
ZhengHui Hu†, Julie Rousseau‡, Nalini Venkatasubramanian*
*Dept. of Computer Science, University of California, Irvine, USA, {mkenne, nalini}@uci.edu
**Div. of Computer Science & Mathematics, Oak Ridge National Laboratory, datepa@ornl.gov
†ImageCat, Inc, Long Beach California, USA, {rte, zh}@imagecatinc.com
‡Div. of Geriatrics & Gerontology, School of Medicine, UCI Health, jroussea@hs.uci.edu

***Abstract*—Efficient resource allocation is crucial in many domains, particularly in senior care, where assigning resources to older adults must consider uncertainties associated with vulnerable populations. In collaboration with Senior Health Facilities (SHFs) and domain experts, this paper presents *iFair*, a novel framework designed to assist decision-makers in equitably allocating scarce resources to older adults. *iFair* was prototyped in the context of ongoing work on a data exchange platform, *CAREDEX*, used for enhancing older adults' resilience during disasters. A key novelty of *iFair* focuses on aligning resident preferences with resources in urgent situations, expediting care, and enhancing task efficiency. We integrate static and dynamic environmental data, including facility layouts and sensor data, with detailed resident profiles to cater to the individual needs and preferences of residents. While our framework primarily focuses on allocation within facilities, it also extends to a regional scale to support the planning and transfer of seniors to mutual aid facilities. Our experiments adapt data from a real SHF to emulate resource allocation in an emergency fire evacuation setting and highlight the delicate balance that decision-makers can achieve between efficiency and fairness.**

***Index Terms*—fairness, resource allocation, senior health care, decision-making.**

I. Introduction

Scheduling in human-centric processes, especially in senior care settings, involves aligning limited resources with high demand, a challenge critical for timely care in healthcare and emergency services. Hence, efficient scheduling is key to managing staff shortages in senior health facilities (SHFs), reducing hospital wait times, and ensuring prompt emergency response [1], [2]. Domain experts highlight that senior care resource allocation often targets staffing and staff time distribution [3]. The COVID-19 pandemic highlighted the magnification of equitable resource allocation challenges, leading to a comprehensive review by the Agency for Healthcare Research and Quality (AHRQ) and encouraging new research into shortages of equipment, medication, and staff [4]–[6]. Along with disparities in COVID-19 vaccine distribution, this accentuated the need for equitable resource allocation, especially as mortality rates among seniors soared with limited critical medical resources like ventilators and ICU beds [7], [8]. Conventional resource allocation guidelines and methods may fall short in human-centric situations, especially in scenarios requiring urgent decision-making. For example, in a hypothetical fire evacuation setting, new caregivers may choose whom to evacuate at random, leading to wasted time and effort. However, prior knowledge of residents' needs and preferences could facilitate this process, motivating the need for informed resource allocation strategies in emergencies.

Older adults, including those with disabilities or physical limitations, often face challenges in accessing essential healthcare services due to resource scarcity. However, current resource allocation methods offer broad solutions, lacking the fine-grained personalization needed for optimal care. Hence, there is a need for a novel approach that is fair, considers the complexities of the environment, and the dynamic state of the residents, accounts for their preferences, and adapts to changes in real time. *iFair* incorporates those by leveraging detailed resident data, including their Activities of Daily Living (ADLs), to support decision-makers and stakeholders.

Allocating resources to older adults must involve considerations like individual needs, resource availability, and care urgency, making fair distribution challenging. Decisions must balance fairness and efficiency, significantly impacting lives. Emergency situations add pressure for swift decision-making, complicating the assurance of optimal choices. At the regional level, resource allocation becomes even more complex, especially during evacuations from disasters. Various factors must be considered, including in-facility conditions, resident relocation, resource availability at mutual aid facilities, real-time traffic, and other special needs. Our experiences with wildfires and earthquakes in California have informed our approach, detailed in our discussion on using relocation tools for regional allocation in - §9. Key contributions include:

- A mathematical model to capture a digital twin representation of a SHF, i.e. facility, residents, resources, etc - §III-A.
- A new definition of fairness for equitable resource allocation which considers human needs and preferences - §III-B.
- A novel framework for fair senior care resource allocation that integrates expert insights and our experience with SHFs. Under this framework, we propose a task-based assessment method for scoring and prioritizing older adults - §IV.
- A detailed evaluation of our framework using an emulated-based approach with data adapted from a real SHF - §V.

2693-8340/24/$31.00 ©2024 IEEE
DOI 10.1109/SMARTCOMP61445.2024.00025

II. Related Work and Limitations

In this section, we review related literature for efficiency and fairness in resource allocation methods and highlight the limitations of such approaches in addressing the specific needs of vulnerable, dynamically changing populations.

Resource allocation has been studied in many settings such as economics [9], education [10], and transportation [1]. These works primarily focus on the *efficiency* of allocation and optimize metrics for completion time [11] and resource usage [12] under varying conditions. In healthcare, resource allocation typically consists of assigning human (e.g., nurses, doctors) and equipment (e.g., beds, wheelchairs) resources to residents and patients. In the context of hospital systems, this has led to metrics such as hospital bed occupancy rates (BOR) and patient length of stay (LOS) [13]. Early work in this context has leveraged solutions from other optimization problems, including makespan [11], MULTIFIT [14], and TSP [15]. These solutions rely on heuristics [14] and evolutionary algorithms [16] for their speed of computation and simplicity, as well as mixed integer linear programs [13] to optimize BOR and LOS. However, for senior health care, these approaches struggle to accurately model complex/dynamic interdependencies (e.g., changing and long-term medical conditions of patients) and the needs of older adults without requiring extensive constraints. The interdependencies between specific resources (a specific caregiver) and a resident influence the feasibility and efficacy of tasks. Recent efforts have explored queuing theory [17] and Markov decision processes [1], [2] to address different system dynamics; however, they lack consideration for preferences and the multifaceted needs of patients and may become complex.

When resource allocation involves real people, it becomes increasingly vital to consider the issue of *fairness* and its trade-offs with efficiency [18]. Several definitions for the concept of fairness in healthcare have emerged in literature, characterized by the distribution of resources to residents and the degree to which they satisfy individual health/well-being needs. For instance, proportional fairness [19] aims to allocate resources proportional to each resident's needs, while max-min fairness [20], [21] looks to maximize the minimum resource share among residents. Other efforts have considered more complex settings with multiple resources and heterogeneous requests [22], as well as sets of desirable properties, e.g., envy-freeness [12] and pareto-efficiency [19], [22]. A recent study [23] also identified several key criteria for prioritizing residents/patients and considers how sickness, prognosis, and waiting lists should be utilized while allocating resources to minimize waiting/evacuation time. In our work, we leverage max-min fairness to allow equal opportunities for resource access based on these needs.

In dynamic environments, e.g., emergencies, the resource allocation problem must consider the impact of time alongside efficiency and fairness [22], [24]. Traditional approaches rely on scheduling strategies such as shortest job first [25], round robin [26], and priority-based schemes [23], [26], [27]. However, these methods discard considerations of preference and the specific needs of residents and patients. Efforts for a more realistic representation of healthcare settings have considered bottlenecks in resources [28] and patient urgency [29], but rely on manual appointment scheduling for a larger timescale. For emergencies that run under shorter timescales, such scheduling becomes impractical. To address this, iFair presents a customized strategy that considers the inherent uncertainties in an emergency, along with the dynamic states of vulnerable populations, including their distinct needs, preferences, and immediate circumstances.

III. Problem Formulation

A. Modeling Key Components

This paper tackles the challenge of fair resource allocation for older adults in senior care facilities. We model the static and dynamic elements of a senior care *facility*, and the *resources* and *tasks* requiring allocation to *residents*. Then, we characterize the *event* and *allocation* processes.

*1) **Facility**:* We model a facility as a graph $\mathcal{G} = (V, E)$ where nodes $v_i \in V$ represent different areas (e.g., rooms, nurse's office) and weighted edges $e_{ij} \in E$ denote a pathway between v_i and v_j with travel time ω_{ij}. We denote two key areas in the facility: exits v_{Exit} and an evacuation area v_{Evac}. The dynamic state of a node v_i is captured by $state(v_i, t)$ representing the condition of v_i at time t, which can be one of $\{Impacted, Clear, Unknown\}$. This is used to pinpoint emergency sites and their proximity to residents. We also model the traversal delays using $access(v_i, t)$, which could occur in an emergency (e.g., due to smoke, or narrow exits).

*2) **Resources**:* We consider two types of resources: *humans* such as nurses and firefighters, and *equipment* such as medication and wheelchairs. Human resources $h_i \in \mathcal{H}$ are characterized by a headcount $\mathcal{C}_i$ and skill-set $\mathcal{S}_i$. We model $\mathcal{S}_i$ as a tuple of personal attributes (e.g., gender) and skills defined in the O*NET resource center skill ontology [30], which includes certifications (e.g., basic/critical care) and abilities (e.g., mobility aid). Equipment resources, $e_i \in \mathcal{E}$ are characterized by their quantity q_i, reusability r_i and type γ_i. In our context, we define reusability to distinguish between single-use items (e.g., syringes, medication) and multi-use items (e.g., portable tanks). The type of equipment resource defines its purpose and capabilities, e.g., a wheelchair is used for mobility and can be manual or electric. For our allocation problem, we also consider the real-time location $loc(\cdot, t)$ and status $avail(\cdot, t)$ of human and equipment resources.

*3) **Tasks**:* We support a set of tasks $\tau_i \in \mathcal{T}$ with dependencies denoted by $\mathcal{T}_D$, such as in Fig. 1. Tasks range from relocating resident or resource entities (which we denote abstractly as $\mathcal{X}$) to specific facility locations $move\_to(v_i)$, to managing entities via $get\_entity()$ and $release\_entity()$. The core tasks, color-coded in blue in Fig. 1, involve providing basic and/or critical care. Denoted as $provide\_basic\_care()$ and $provide\_critical\_care()$, they range from ADLs such as dressing and toileting to more advanced medical needs.

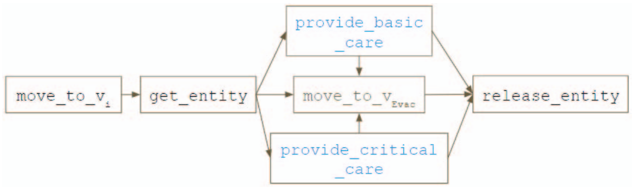

Fig. 1. Task Dependency Graph.

Each task is defined by its name $\mathcal{N}$, entities involved $\mathcal{X}$, and the resources needed $Req(\mathcal{E}, \mathcal{H})$. We model the changing conditions of an emergency through task updates $task\_update(\tau_i, \mathcal{X}, t)$. We define the overall task completion time in Eqn. 1, which encompasses the path traversal times ω_e, delays for space access $access(\cdot)$, and additional time incurred for allocating a resident's non-preferred resources λ_p (described later). To optimize task execution, Dijkstra's shortest path algorithm is employed.

$$TTC() = \sum_e \omega_e + \sum_v access(v_i, t) + \lambda_p, \forall_{e,v} \in \mathcal{G} \quad (1)$$

*4) **Residents**:* Let $P_i \in \mathcal{P}$ denote residents in the senior care facility. Each resident is characterized by a 3-tuple $\{\mathcal{B}, \mathcal{D}, \mathcal{F}\}$, denoting basic attributes (e.g., name, gender, age, DOB, room), their diagnoses (e.g., dementia, arthritis, obesity, etc.), and preferences for specific resources. In our model, preferences are quantified on a numerical scale ranging from -5 (strong dislike) to $+5$ (strong preference). Table I indicates these preferences for basic care tasks. This scaling allows us to measure the impact on task efficiency and residents' satisfaction. Positive values indicate resources that enhance operational performance, while negative values increase task completion time; directly reflecting the operational costs of utilizing less favored resources.

A resident's satisfaction level $\mathcal{LS}()$ is the sum of their preferences $\mathcal{F}_i$, for all resources used in a task τ_i, i.e., $\mathcal{LS}() = \sum_{i=0}^{n_{P_i}} \mathcal{F}_i$, where n_{P_i} is the number of resources allocated to resident P_i. The extra time λ_p, for using less preferred resources is computed from summing negative $\mathcal{F}_i$ values, i.e., $\lambda_p = -\sum \alpha_i \cdot \min(0, \mathcal{F}_i)$. α_i quantifies how much a preference impacts task efficiency. For an emergency scenario, let the criticality score C_{score}, assess the urgency of residents' needs based on a set of important attributes $\mathcal{I}$ with corresponding weights ω_a. We track changes in resident conditions (e.g., head injury from falling) using a health report form $\mathcal{HRF}$ and adjust attribute weights and necessary resources which are described later in §IV-A. A resident's location is captured using $loc(P_i, t)$; we assume that residents have expected locations at specific times, as depicted in Fig. 2.

Fig. 2. Resident Expected Location.

TABLE I
RESIDENT PREFERENCES FOR BASIC CARE.

	Joe	Mary	Bob	Carol	Don
h_1	+5	-5	0	-4	0
h_2	+2	-3	0	-1	0
h_3	-5	+3	0	+4	0
e_1	0	0	0	+5	0
e_2	0	-1	0	-2	-5

Consider a resident, Joe, who requires two-person assistance and equipment e_2 (See Table I). If he was assigned resources h_2, h_3, and e_2 under impact rates α_{h_2}=3, α_{h_3}=2, and α_{e_2}=4 (in minutes per $\mathcal{F}_i$), we find that λ_{Joe}=10 extra minutes, and $\mathcal{LS}(Joe)$=-3, suggesting slight dissatisfaction with his allocated resources.

*5) **Event and Allocation**:* Our analysis prioritizes urgent events like fires that require immediate action. An *event* is defined by its starting point, with $v_o = loc(event, t{=}0)$, and its propagation, $prop(v_i, v_j)$ within the facility, which increases edge weights ω near the affected nodes. Allocations $\mathcal{A}$ are denoted by $\{P_i, h_i \,||\, e_i, t_s, t_e\}$, indicating that a resident P_i is allocated a resource h_i or e_i during time period (t_s, t_e). Notably, resources assigned to the same task might start simultaneously but end at different times, such as a nurse and a consumable (medication). We also note that the human resource count in senior facilities is much lower than that of residents, i.e., $\sum_{i=1}^{|\mathcal{H}|} \mathcal{C}_i << \sum_{j=1}^{|\mathcal{P}|} P_j$.

B. Problem Statement

In our model, fairness is defined as the balance between prioritizing urgent resident needs, reflected by criticality scores C_{score}, and achieving equitable satisfaction levels $\mathcal{LS}()$ across residents. This dual objective ensures resources are initially allocated to those in immediate need based on their conditions and needs while striving to meet the satisfaction of all residents uniformly. Yet, emergencies introduce a balance challenge between fairness and efficiency; for instance, during a fire, swift evacuation might override individual preferences for everyone's safety. Hence, fairness is prioritized without compromising urgent safety measures, leading to the formulation of the fair resource allocation problem as shown in Eqn. 2.

$$\begin{aligned}
\min \quad & \sum(\omega_1 \cdot \lambda_p - \omega_2 \cdot \mathcal{LS}()) && \text{(2a)}\\
\text{subject to} \quad & \lambda_p = \sum |\mathcal{F}_i| \times \alpha_i, \forall_{\mathcal{F}_i} < 0 && \text{(2b)}\\
& \mathcal{LS}() = \sum_{i=0}^{n_{P_i}} \mathcal{F}_i && \text{(2c)}\\
& \omega_1 + \omega_2 = 1 \ , \ 0 \leq \omega_1, \omega_2 \leq 1 && \text{(2d)}
\end{aligned}$$

This corresponds to a multi-objective optimization problem with two objectives where one consists of accomplishing efficiency by minimizing $\sum \mathcal{F}_p$, incurred when assigning non-preferred resources to residents and achieving fairness by maximizing $\sum \mathcal{LS}()$, the residents' satisfaction. ω_1 and ω_2 are the efficiency and fairness weights respectively. For instance, setting $\omega_1 = 0.4$ and $\omega_2 = 0.6$ indicates a greater emphasis on fairness over efficiency in the allocation.

***Theorem**: Our fair resource allocation problem is NP-Hard.*
Proof: We demonstrate its NP-Hardness by reducing from the makespan optimization problem, a classic NP-Hard problem

in scheduling theory. Consider an instance of the makespan problem with n jobs and m machines, where each job i has a processing time p_{ij} on machine j. We construct an analogous instance of our resource allocation problem as follows; each job corresponds to a resident, and each machine to a resource (either human or equipment). Let $p_{ij} = TTC()$ represent the processing time of resident i when assigned to resource j. In the simple case of identical machine scheduling problems where the processing time of each job is the same on each machine, minimizing the weighted average completion time is NP-hard by reduction from the Knapsack problem [31]. Even if there are only two machines, the problem remains NP-Hard by reduction from the partition problem [32]. In our resource allocation problem, the processing times are unrelated. For instance, machine i could have $p_{ij} > p_{ij'}$, while machine i' could have $p_{i'j} < p_{i'j'}$. This corresponds to an instance of the unrelated parallel machines scheduling problem where the objective is to minimize the makespan (Eqn. 3).

$$\min \quad \sum_{j=1}^{n} p_{ij} x_{ij} \leq C_{max} \tag{3a}$$

$$\text{subject to} \quad \sum_{i=1} x_{ij} = 1, \forall_{ij} \in \{0, 1\} \tag{3b}$$

Our reduction from the well-established NP-Hard makespan problem shows that our resource allocation problem is at least as complex as the makespan problem. Consequently, our problem inherits the NP-Hard classification. In the subsequent section, we outline our methodology for achieving fairness.

IV. iFair Approach

A. Overview of iFair Approach

We prototyped the iFair framework in the context of CAREDEX: a data exchange platform that aims to enhance the resilience of older adults in aging communities. We assume the availability of static healthcare and dynamic smart space information about residents; sensors embedded in the space are used to collect data about residents' location and movement [33]–[35]. Information from health records and caregiver logs can provide additional data about health and ADLs to create a comprehensive picture of each individual's needs. This information is then shared with other stakeholders including first responders, family, and caregivers [36]. Fig. 3, provides an overview of key components in the iFair framework.

(a) Physical Infrastructure and Occupants: First, we collect data about the residents and resources in the facility using a variety of records and sensors for localization and tracking, such as motion sensors, emergency call buttons, and Wi-Fi sensing. Note: Indoor localization is not the main focus of this paper; however, through the CAREDEX project, we have conducted multiple drills that allow us to approximate the location of participants. Thus, we have some real-world experience in this area. All data collection complies with the Health Insurance Portability and Accountability Act of 1996 (HIPAA) regulations, utilizing encryption and data anonymization to ensure the privacy and security of resident information [34] (out of scope for this paper).

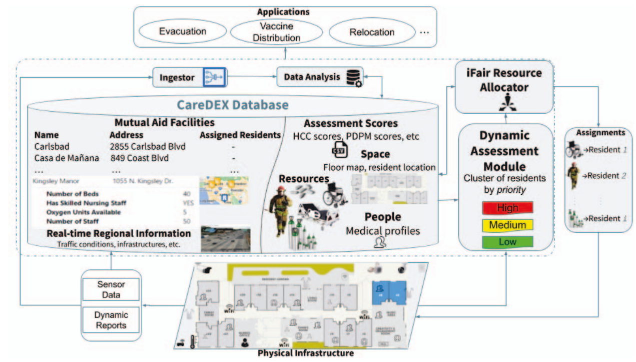

Fig. 3. Overview of the iFair Framework.

(b) Resident/Resource Database: Next, we store and analyze this data to gain insights about the residents. We also use the database for storing the residents' medical profiles and ICD-10-CM (International Classification of Diseases, Tenth Revision, Clinical Modification) [37] coding data including real-time regional information about mutual aid facilities (e.g.: number of beds or oxygen units available).

(c) Dynamic Assessment Module: We use this module to assess the residents and update their criticality scores, a measure of their wellness and needs. Furthermore, we cluster the residents into three groups based on their criticality. This module also handles real-time changes that affect the priority of residents. e.g.: if a resident classified as low-priority based on criticality is injured during an evacuation, we reassess them and update their criticality score.

(d) Resource Allocator: Based on all known information, this module assigns the resources (see Alg. 3). Overall, this framework can serve a range of purposes such as facilitating the evacuation of residents, improving the distribution of vaccines during an outbreak, and enabling swift relocation of residents in the face of a natural disaster.

B. Task-Based Assessment and Scoring

We collaborated with geriatric experts at the University of California Irvine (UCI) to explore leading assessment tools for evaluating older adults. This exercise led to the adoption of the ICD-10-CM system for resident assessment and scoring. The ICD-10-CM, developed by the Centers for Disease Control and Prevention (CDC) and based on the World Health Organization's (WHO) alphanumeric disease classification codes, is widely used in the U.S. healthcare system for disease classification. In line with HIPAA regulations, healthcare professionals use these codes for patient assessments, assigning risk scores for various purposes such as determining care levels, classifying healthcare services and pricing, and assessing cognitive and mobility functions among older adults. Assessments are performed when a resident is admitted to a facility, quarterly, and annually. In addition, residents are assessed whenever they experience a significant change in status, and whenever the facility identifies a significant error in a prior assessment.

Algorithm 1: Offline Criticality Score Assessment

Input: Residents $\mathcal{P}$, Important attributes $\mathcal{I}$, Desired attribute value v_α, Weight of attribute ω_α
Output: List of Criticality Scores

1 $List\_C_{score} \leftarrow \varnothing$
2 **for** $resident\ in\ \mathcal{P}$ **do**
3 $\quad C_{score} \leftarrow 0$ `// to update using ICD-10-CM`
4 $\quad$ **for** $\mathcal{D}_i \leftarrow resident.diagnoses$ **do**
5 $\quad\quad C_{score} \leftarrow C_{score} + hcc\_val(\mathcal{D}_i)$
6 $\quad C_{score} \leftarrow \sum_{\alpha \in \mathcal{I}} \omega_\alpha \cdot HCC\_value_\alpha, \forall_{\alpha = v_\alpha}$
7 $\quad$ Add C_{score} to $List\_C_{score}$
8 **return** SortDescending($List\_C_{score}$)

a) Offline Criticality Score - (computed apriori based on medical profile): For evaluating residents' physical conditions, we analyze attributes from their medical profiles, employing a risk-adjustment model to calculate their Risk Adjustment Factor (RAF) scores. This study adopts the Hierarchical Condition Category (HCC) model [38] for this purpose.

Consider a resident with a severe head injury coded as T31.11, equating to an HCC value of 0.486. Adding dementia increases their score by 0.346. A BMI between 40.0-44.9 adds another HCC value of 0.25. Summing these values gives $C_{score} = 0.486 + 0.346 + 0.25 = 1.082$ (refer to lines 4-5 of Alg. 1). This method is applied to all residents, noting that those with a score over 1.0 are regarded as having severe health issues. During an urgent evacuation, attributes like the weight of a resident and their ability to move independently are crucial. To address this, we define important attributes $\mathcal{I}$ and desired values v_α, e.g., $\mathcal{I} = \{BMI, Ambulatory\}$ with $v_{BMI} > 25$ for overweight, and $v_{Ambulatory} = False$ for mobility issues. A set of weights is then applied to emphasize these attributes' importance. Consequently, residents matching these criteria have their HCC values adjusted, thereby updating their criticality score as outlined in line 6 of Alg. 1. This adjustment yields a weighted C_{score}, our measure for offline criticality.

b) Online Criticality Score - (computed in real-time based on event urgency): During emergencies, criticality scores are updated to reflect the urgency of care needed by residents, as outlined in Alg. 2. In a fire scenario, for instance, a resident's criticality score is modified based on their closeness to the fire, enabling a reevaluation of their priority. This could lead to assigning a higher priority to those in immediate danger, thus allowing for a dynamic response to the unfolding situation. By combining the CDC's established medical scoring framework with real-time data from sensors across the facility, we dynamically gauge the new criticality of each resident.

Algorithm 2: Online Criticality Score

Input: $\mathcal{P}$, $\mathcal{G}$, $List\_C_{score}$
Output: $List\_C_{score}$

$v_o \leftarrow \mathcal{G}.loc(event, t = 0)$
for $P_i \leftarrow \mathcal{P}$ **do**
$\quad d_{P_i} \leftarrow ShortestDistance(loc(P_i, t), v_o)$
$\quad U_{P_i} \leftarrow \frac{1}{d_{P_i}}$ `// compute urgency`
$\quad U\_List.add(U_{P_i})$
for $U_{P_i} \leftarrow U\_List$ **do**
$\quad$ Normalized $U'_{P_i} = \frac{U_{P_i} - min(U\_List)}{max(U\_List) - min(U\_List)}$
$\quad$ Update $List\_C_{score}[P_i] \leftarrow List\_C_{score}[P_i] * U'_{P_i}$
return SortDescending($List\_C_{score}$)

After assessment and scoring, we allocate resources to residents as shown in Alg. 3, cycling through sorted criticality scores to match residents with tasks, which involves setting task names (e.g., $provide\_basic\_care()$), identifying dependencies, and determining necessary resources with preferences (lines 3-6). We assign preferred resources within d_{max} and to avoid starvation, any human available resource (lines 7-12). Then compute task completion times and satisfaction (lines 13-14), updating resource states and storing these in the allocation record (lines 15-21). Additionally, we maintain a health report form ($\mathcal{HRF}$) where we include residents experiencing changes in health conditions and use it to update their criticality scores.

Algorithm 3: Resource Allocation

Input: Online criticality scores $List\_C_{score}$, Distance threshold d_{max}, Health report form $\mathcal{HRF}$
Output: Allocation records $\mathcal{A}_R$

1 $\mathcal{A} \leftarrow \emptyset; \mathcal{A}_R \leftarrow \emptyset$
2 $t_s \leftarrow$ GetCurrentTime()
3 **while** $\neg$ *isEmpty*($List\_C_{score}$) **do**
4 $\quad C_{score} \leftarrow \max(List\_C_{score})$
5 $\quad P_i \leftarrow$ GetCorrespondingResident(C_{score})
6 $\quad R_{eq}, \mathcal{F}_i \leftarrow$ GetTaskDetailsAndPrefs(P_i)
7 $\quad$ **for** $rr \in R_{eq}$ **do**
8 $\quad\quad r \leftarrow$ GetResource($\mathcal{F}_i, max\_dist = d_{max}$)
9 $\quad\quad$ **if** $r ==$ **null** `// Assign human`
10 $\quad\quad$ **then** $r \leftarrow$ NextAvailableHuman() ;
11 $\quad\quad \mathcal{A} \leftarrow \mathcal{A} \cup \{(P_i, r, t_s)\}$
12 $\quad\quad$ SetResourceState(r,"in-use")
$\quad$ `// Get Task completion time, satisfaction`
13 $\quad TTC() \leftarrow \sum \omega_e + \sum access(v_i, t_s) + \lambda_p$
14 $\quad \mathcal{LS}() \leftarrow \sum_{i=0}^{n_{P_i}} \mathcal{F}_i$
15 $\quad$ **for** r *in* $\mathcal{A}$ **do**
16 $\quad\quad$ SetResourceState(r,"unavailable")
17 $\quad\quad$ **if** *reusable*(r) **then**
18 $\quad\quad\quad t_s, t_e \leftarrow t_e, t_e + TTC()$
19 $\quad\quad\quad$ SetResourceState(r,"available")
20 $\quad\quad\quad \mathcal{A} \leftarrow \mathcal{A} \cup \{(P_i, r, t_e)\}$
21 $\quad \mathcal{A}_R \leftarrow \mathcal{A}_R \cup \{(TTC(), \mathcal{LS}(), \mathcal{A})\}$
22 $\quad$ **if** $\neg$*isEmpty*($\mathcal{HRF}$) **then**
23 $\quad\quad$ **for** $i, sc \leftarrow$ *GetUpdatedCriticalityScores*($\mathcal{HRF}$) **do**
$\quad\quad\quad List\_C_{score}[i] \leftarrow sc$;
24 **return** $\mathcal{A}_R$

V. Experimental Evaluation

Experimental Setup: We evaluated our framework through emulation, utilizing anonymized data modeled after a real-world senior care facility in Orange County, USA, including 30 elderly residents' profiles (age, gender, medical information). With input from experts in the Department of Family

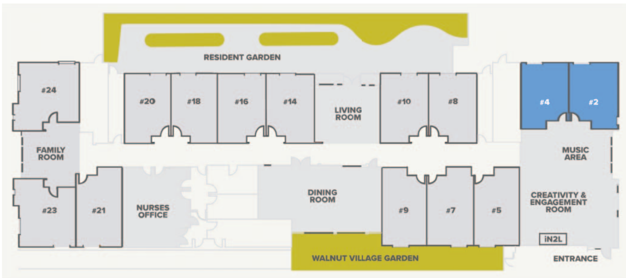

Fig. 4. Floor Plan of Small Facility.

Medicine, Division of Geriatric Medicine and Gerontology at UCI Health, we generated synthetic data for 400 resident profiles to facilitate larger population experiments. The offline criticality score was evaluated using the CASAS dataset [39], featuring real-world data on Activities of Daily Living (ADL), activity scores, and diagnoses for 400 residents from Washington State University. We mapped ICD-10-CM diagnoses to compute scores using HCC coding, based on residents' medical information. A simulator (written in Python) was developed to emulate resource allocation within a facility, incorporating actual floor plans (Fig. 4) from a partner facility.

Experimental Plan: When a disaster strikes, older adults are disproportionately affected [40]. During this type of event, the human resources (i.e. medical staff, first responders, etc.) needed to evacuate this frail population from their SHF to a safe location are minimal. Hence, we crafted a use case where rapid evacuation is imperative. We map this scenario to a parallel machine scheduling problem translating staff to machines and residents to jobs. Using the formulation in Section III, our goal is to reduce evacuation times (i.e. optimize the makespan C_{max} [41]) as well as the residents' satisfaction. Accordingly, the key metrics we utilize in our evaluation are time-to-completion (an efficiency/latency metric) and level of satisfaction (a fairness metric).

We next discuss strategies that are used as baselines and comparison points for iFair. Senior care facilities have emergency plans in place that correspond to the following algorithms: Shortest Job First (SJF), which aims to evacuate residents who can be moved with the least effort and in the shortest amount of time. This includes those who are fully ambulatory, require minimal assistance, or are closest to the exits. Longest Job First (LJF), on the other hand, prioritizes the most critical residents, requiring the most time and resources to evacuate. This could include those with mobility issues, those requiring medical equipment, or residents located in the most challenging parts of the facility.

To ensure a fair distribution of tasks among staff members, we also developed a Clustered Round Robin (C-RR) strategy. Here, each staff is assigned a fixed number/group of residents. For example, staff and equipment may be assigned to specific units/floors for a specific period to accomplish the task before moving on to the subsequent floor. We also developed a Criticality-aware Clustered Round Robin (CC-RR) strategy

TABLE II
COMPARATIVE OUTCOME WHEN EVACUATING RESIDENTS WITH PREFERENCES (WP) VS. WITHOUT PREFERENCES (NP).

Priority	Residents	C_{score}	$\mathcal{LS}$-WP	TTC-WP	$\mathcal{LS}$-NP	TTC-NP	λ_p
HIGH	Joe	2.741	5	10	4	10	0
	Mary	2.634	5	12	2	27	15
	Bob	2.527	6	6	-6	6	0
	...	...	...	...	...	...	...
MED	Don	1.813	1	10	-6	18	8
	Georgia	1.806	1	12	-6	22	10
	Laura	1.799	1	14	-6	18	4
	...	...	...	...	...	...	...
LOW	Mona	0.708	-1	17	2	18	1
	Gennifer	0.601	0	18	2	20	2
	Karole	0.0	5	11	-4	15	4
	...	...	...	...	...	...	...

which is an improved version of the C-RR technique. CC-RR consists of assigning resources first to high-priority residents, then when there are no more residents in that cluster, it proceeds as C-RR alternating between medium-priority and low-priority residents for a given amount of time. Let us note that strict adherence to any single plan could lead to liability issues based on negligence [42]. LJF policies may overlook urgent needs for less critical, resource-intensive cases. Conversely, SJF risks neglecting high-need residents for faster evacuations, while RR's fixed order may delay aid to urgent cases.

We use these algorithms to assess the effectiveness and efficiency of our framework across various facility sizes. By implementing these in both smaller settings, where resource constraints may differ, and larger institutions with more complex logistical challenges, we comprehensively evaluate our approach. This enables us to adapt and refine our strategies, optimizing evacuation procedures and ensuring equitable care delivery in diverse operational environments.

Experimental Results: Although the evacuation time for some residents, like Joe, remains the same (10 minutes) whether preferences are used or not, indicating that his evacuation is not impacted by any additional time due to preference settings, most evacuation times observed in scenarios where resident preferences were considered (TTC-WP) compared to those where these were not included (TTC-NP), reveal significant differences, as quantified in Table II. Notably, Mary experiences a 15-minute increase in evacuation time without preferences, rising from 12 to 27 minutes. This indicates that ignoring preferences can lead to inefficient routing or resource allocation, significantly increasing evacuation times.

Moreover, computing λ_p across different resident profiles further illuminates this disparity. This measure quantifies the additional time or delay potentially incurred when resident-specific preferences and needs are overlooked. For instance, Mary's high λ_p value of 15 suggests considerable neglect of her specific needs or location, which is adequately addressed when preferences are considered, resulting in a more optimized and faster evacuation process under the TTC-WP scenario.

iFair, particularly in a small facility as shown in Fig. 5, achieves the lowest variability in satisfaction, with standard deviations of 0.58 for high, 2.12 for medium, and 2.47 for low priority levels, indicating more equitable satisfaction

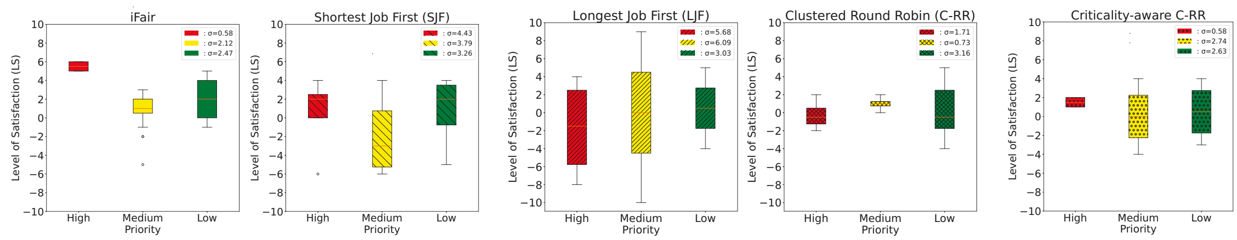

Fig. 5. Residents' Satisfaction Across Different Approaches for a Small Facility

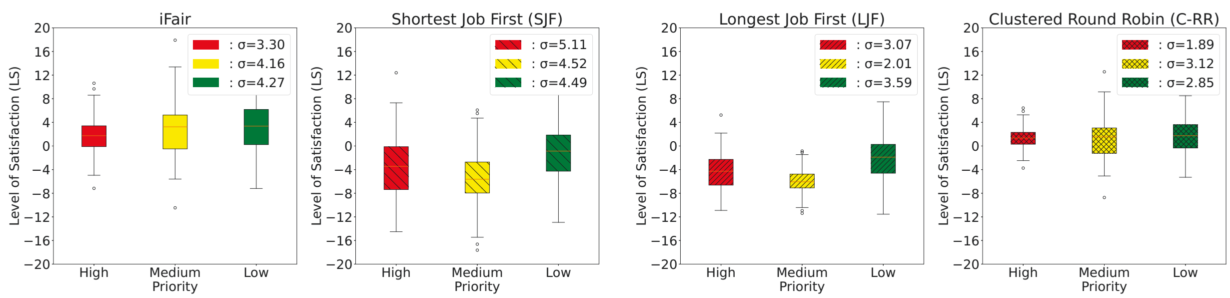

Fig. 6. Residents' Satisfaction Across Different Approaches for a Large Facility

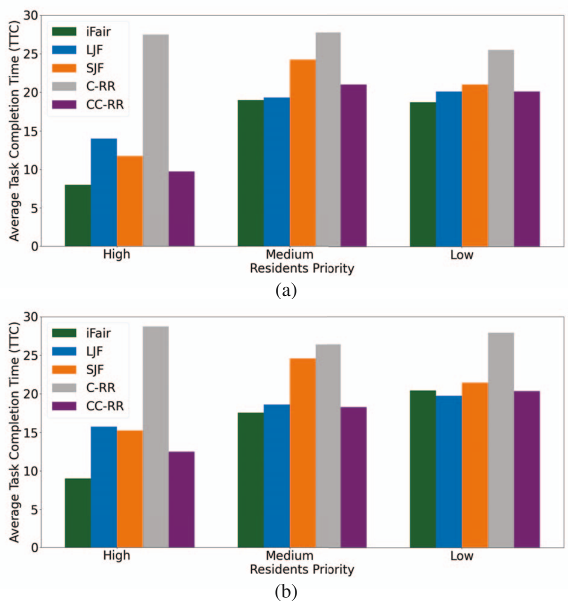

Fig. 7. Comparing Average Task Completion Times: (a) Small Facility vs. (b) Large Facility.

among residents. It surpasses the other methods in terms of satisfaction spread. The Clustered Round Robin method shows moderate variability, suggesting the importance of a more personalized approach.

In larger facilities, as illustrated in Fig. 6, iFair maintains lower variability in satisfaction, highlighting its effectiveness and scalability in evacuations that consider resident preferences. The figure shows that iFair outperforms other methods by achieving higher satisfaction levels across high, medium, and low-priority residents. The satisfaction levels depicted in Fig. 6 for iFair are higher compared to SJF and LJF, especially for residents of medium and low priority, where the need to balance resources can often leave these groups less attended in other models. Even when compared to Clustered Round Robin (C-RR), iFair demonstrates a more favorable satisfaction profile. In Fig. 7, iFair consistently provides the shortest average task completion time (TTC) for high-priority residents in both facilities, demonstrating efficient prioritization. LJF and SJF show moderate TTCs, indicating less effective prioritization, while C-RR records the longest TTC, showing a slower response. Likewise, the efficiency gap between iFair and the other methods, especially for high-priority tasks in large facilities, highlights iFair's scalability and effective prioritization. CC-RR, despite high TTCs, suggests potential scalability close to iFair's performance.

The Gantt charts in Fig. 8 indicate that iFair achieves a 68-minute makespan, setting a benchmark. Compared to other methods, it improves the makespan by 20% over LJF, 35.85% over SJF, 22.47% over CC-RR, and 37.04% over C-RR. iFair's efficiency stems from prioritizing urgency and preferences, contrasting with methods leading to longer completion times. Overall, our findings demonstrate that our framework is equitable, as supported by the satisfaction of the residents, and it is also an efficient system.

VI. Extensions for Regional Resource Allocation and Decision Support

Planning and coordinating resources for distressed seniors is vital within and across facilities. We explore regional awareness in developing cross-facility resource-sharing strategies for elderly care, including senior relocation during evacuations.

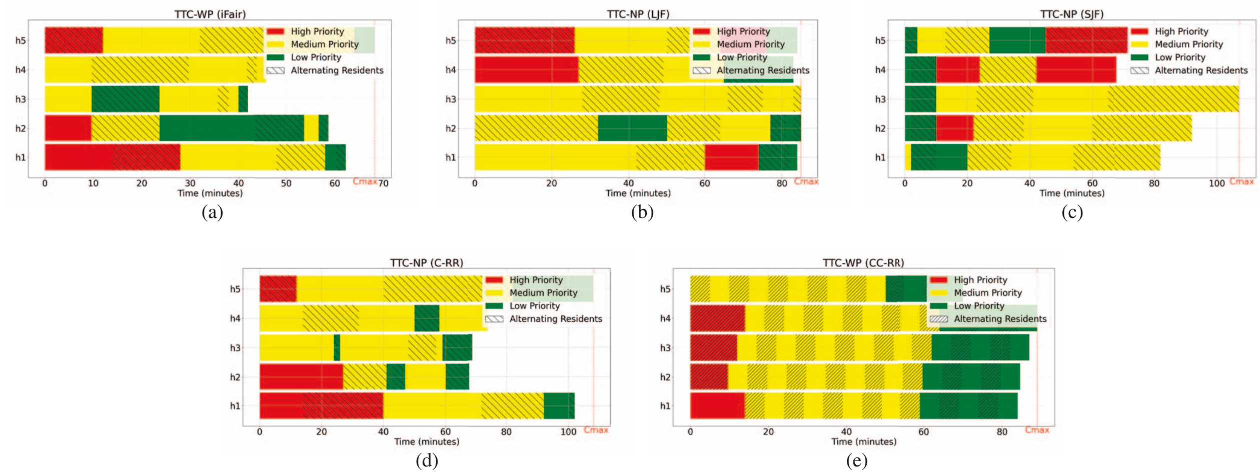

Fig. 8. Gantt Charts of Human Resource Allocation, using (a) iFair; (b) Longest Job First; (c) Shortest Job First; (d) Clustered Round Robin; and (e) Criticality-aware Clustered Round Robin for a Small Facility.

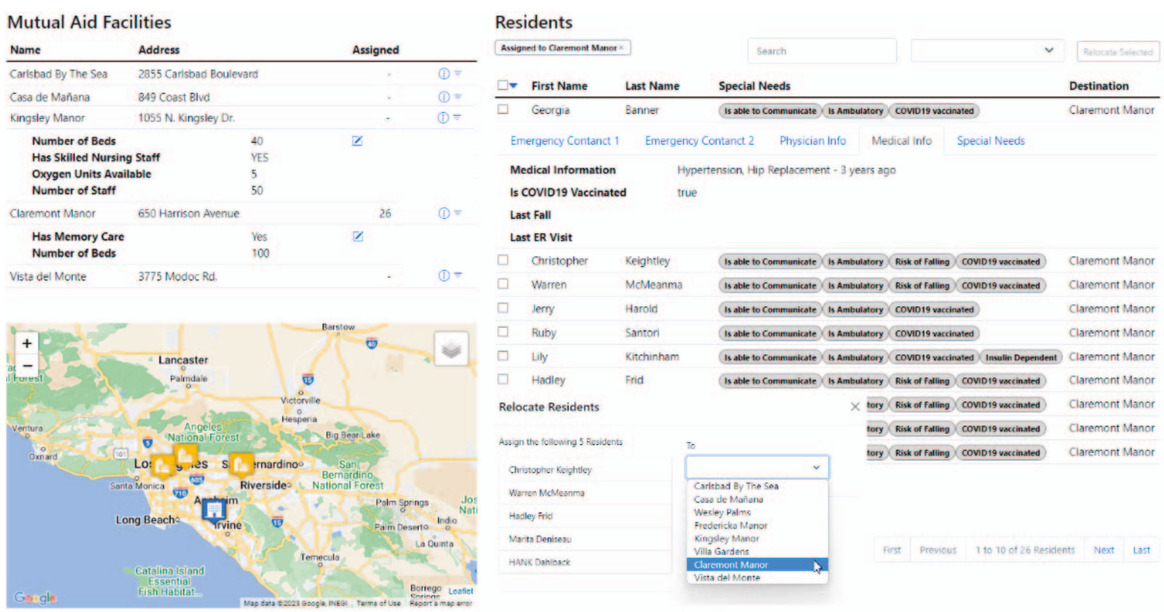

Fig. 9. Relocation Planning Tool

This involves CAREDEX, a dynamic data exchange platform to share regional resource availability (space/beds, caregivers, special equipment such as oxygen, dialyzer) and personalized care needs of relocating individuals for new providers unfamiliar with their unique requirements. CAREDEX features an interactive visualization tool (Fig. 9) for viewing regional impacts and a relocation module for coordinating resources during disasters, offering non-technical decision-makers disaster-specific information such as location, proximity, and severity, and employing simulation tools like FEMA Hazus for estimating damages. Its frontend is developed using Vue.js and incorporates Leaflet and Google Maps for mapping, while the backend utilizes Django REST and stores data in a PostgreSQL database with GIS capabilities. The backend actively collects data from various public hazard sources specific to the hazard type. For earthquakes, it utilizes United States Geological Survey data [43], including Shakemap for immediate impact assessment. A HAZUS-based engine estimates building damage and potential injuries or fatalities. Wildfire information is provided by the National Interagency Fire Center [44], assessing fire proximity to facilities. Additionally, the risk to older adults from wildfires extends beyond physical harm to increased respiratory risk from smoke exposure, highlighted by epidemiological research [45].

NOAA's Hazard Mapping System Fire and Smoke Product [46] provides daily smoke plume data, and AIRNOW [47] offers hourly updates on ozone, PM1.0, and PM2.5 levels, aiding in relocation decisions. We used CAREDEX's regional awareness tool in drill exercises to understand the smoke spread and air quality impacts from nearby wildfires on facilities. For earthquakes, it evaluates impacts on senior care facilities and evacuation sites. CAREDEX integrates data on healthcare facilities' conditions, like local trauma centers, to assess damage at potential relocation sites. It pre-identifies residents' special needs, matching them with mutual aid facilities' resources, and offering relocation options.

VII. Conclusion And Future Work

In summary, this paper highlights the benefits of integrating residents' preferences and real-time criticality assessments in resource allocation for senior care facilities. Leveraging insights from domain experts and our experience with SHFs, we proposed a method to assess and score older adults based on their personal information. Our findings indicate that integrating preferences into allocation algorithms enhances user satisfaction and operational efficiency, ensuring a balanced approach to resource distribution. Beyond emergency evacuations, the iFair framework is adaptable to non-evacuation scenarios such as equipment or staff shortages resulting from unforeseen incidents or disruptions that impact the normal flow of activities in senior care facilities. This adaptability enables iFair to prioritize needs and allocate resources effectively, ensuring continuity of care even when mobility is not required. Such flexibility not only strengthens facility-level operations but also sets the stage for broader applications of this framework. Hence, future work will focus on enhancing resource allocation at the regional scale to facilitate the sharing of resources across multiple facilities.

Acknowledgements

The authors thank Andrew Chio, Fangqi Liu, Prof. Shangping Ren, and members of the Distributed Systems Middleware (DSM) Group at UCI for their valuable input to this work. This research is also supported by DARPA under Agreement No. FA8750-16-2-0021, by the U.S. NSF Grants No. 2044107, 2133391, 2008993, 2245372, and the ECR Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the US Government, DARPA, or the National Science Foundation.

References

[1] G. Pettet *et al.*, "Hierarchical planning for dynamic resource allocation in smart and connected communities," *ACM Transactions on Cyber-Physical Systems*, vol. 6, no. 4, pp. 1–26, 2022.
[2] V. A. Knight *et al.*, "Ambulance allocation for maximal survival with heterogeneous outcome measures," *Omega*, vol. 40, no. 6, 2012.
[3] R. A. Anderson *et al.*, "Resource allocation and resident outcomes in nursing homes: Comparisons between the best and worst," *Research in Nursing & Health*, vol. 21, no. 4, 1998.
[4] E. Emanuel *et al.*, "Allocation of scarce medical resources in the time of covid-19," *N Engl J Med*, vol. 10, 2020.
[5] S. Hempel *et al.*, "Resource allocation and pandemic response: An evidence synthesis to inform decision making," 2020.
[6] J. F. Lynch *et al.*, "Scarce resource allocation in a pandemic: A protocol to promote equity, timeliness, and transparency," *Critical Care Explorations*, vol. 3, no. 6, 2021.
[7] N. A. of Sciences Engineering and Medicine, *Framework for Equitable Allocation of COVID-19 Vaccine*. National Academies Press, 2020.
[8] L. Wedlund *et al.*, "New ml model predicts who may benefit most from covid-19 vaccination," *NPJ Digital Medicine*, vol. 4, no. 1, 2021.
[9] C. A. Maritan *et al.*, "Resource allocation and strategy," 2017.
[10] The Education Trust, "Resource allocation reviews: A critical step to school improvement," 2014.
[11] F. Cappadonna *et al.*, "Makespan minimization of unrelated parallel machines with limited human resources," *Procedia Cirp*, vol. 12, 2013.
[12] R. Fagin and J. H. Williams, "A fair carpool scheduling algorithm," *IBM Journal of Research and development*, vol. 27, no. 2, 1983.
[13] L. Luo *et al.*, "A data-driven hybrid three-stage framework for hospital bed allocation," *Comput. Math. Methods Med.*, vol. 2019, 2019.
[14] M. Kurz *et al.*, "Heuristic scheduling of parallel machines with sequence -dependent set-up times," *Int. J. Prod. Res.*, vol. 39, no. 16, 2001.
[15] G. Chryssolouris *et al.*, "On the resources allocation problem," *The International Journal of Production Research*, vol. 30, no. 12, 1992.
[16] E. Vallada *et al.*, "A genetic algorithm for the upm scheduling problem with sequence dependent setup times," *EJOR*, vol. 211, no. 3, 2011.
[17] R. Gayathri, "An enhanced resource allocation algorithm to allocate hospital beds during the covid19 pandemic," in *ICIRCA*. IEEE, 2021.
[18] D. Bertsimas *et al.*, "The price of fairness," *Operations research*, vol. 59, no. 1, 2011.
[19] T. Bonald *et al.*, "A queueing analysis of max-min fairness, proportional fairness and balanced fairness," *Queueing systems*, vol. 53, 2006.
[20] A. Ghodsi *et al.*, "Dominant resource fairness: Fair allocation of multiple resource types." in *Nsdi*, vol. 11, no. 2011, 2011.
[21] W. Ogryczak *et al.*, "Fair optimization and networks: A survey," *Journal of Applied Mathematics*, vol. 2014, 2014.
[22] J. Le Boudec, "Rate adaptation, congestion control and fairness: A tutorial," *on line*, 2008.
[23] M. H. Yousef *et al.*, "The fair allocation of scarce medical resources: a comparative study from jordan," *Frontiers in medicine*, vol. 7, 2021.
[24] T. M. John *et al.*, "First come, first served?" *Ethics*, vol. 130, 2020.
[25] M. Ji *et al.*, "Parallel-machine scheduling of simple linear deteriorating jobs," *Theoretical Computer Science*, vol. 410, no. 38-40, 2009.
[26] A. Abdelhamid *et al.*, "Implementing and measuring the performance of pb, rr and pbrr scheduling algorithms on atmega32a using freertos," in *NILES*. IEEE, 2023.
[27] I. Sattar *et al.*, "Multi-level queue with priority and time sharing for real time scheduling," *IJMSE*, vol. 5, no. 8, 2014.
[28] W. Huang *et al.*, "Dynamic configuration scheduling problem for stochastic medical resources," *JBI*, vol. 80, 2018.
[29] I. B. Vermeulen *et al.*, "Adaptive resource allocation for efficient patient scheduling," *Artificial intelligence in medicine*, vol. 46, no. 1, 2009.
[30] "O-net resource center," https://www.onetcenter.org//.
[31] E. Horowitz and S. Sahni, "Exact and approximate algorithms for scheduling nonidentical processors," *JACM*, vol. 23, no. 2, 1976.
[32] S. K. Sahni, "Algorithms for scheduling independent tasks," *JACM*, vol. 23, no. 1, 1976.
[33] Y. Lin *et al.*, "Locater: cleaning wifi connectivity datasets for semantic localization," *arXiv preprint arXiv:2004.09676*, 2020.
[34] S. Mehrotra *et al.*, "Tippers: Privacy cognizant iot environment," in *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 2016, pp. 1–6.
[35] A. Chio *et al.*, "Smartspec: A framework to generate customizable, semantics-based smart space datasets," *Pervasive and Mobile Computing*, vol. 93, p. 101809, 2023.
[36] "CareDEX Workshop," https://sites.uci.edu/caredex/.
[37] CDC, "ICD-10-CM Tool," https://icd10cmtool.cdc.gov/.
[38] AAFP, "Hierarchical condition category coding," 2024.
[39] J. Diane *et al.*, "Casas: A smart home in a box," *computer*, 2012.
[40] N. Aldrich and W. F. Benson, "Disaster preparedness and the chronic disease needs of vulnerable older adults," *PCD*, 2008.
[41] M. Englert *et al.*, "The power of reordering for online minimum makespan scheduling," *SICOMP*, vol. 43, no. 3, 2014.
[42] J. G. Hodge Jr, "A legal duty to evacuate patients from healthcare facilities in emergencies," *Health Law.*, vol. 25, p. 20, 2012.
[43] USGS, https://www.usgs.gov/programs/earthquake-hazards/earthquakes.
[44] NIFC, "Nifc open data site," https://data-nifc.opendata.arcgis.com//.
[45] US EPA, "Which populations experience greater risks of adverse health effects resulting from wildfire smoke exposure?" 2022.
[46] O. of Satellite and P. Operations, "Hazard mapping system," https://www.ospo.noaa.gov/Products/land/hms.html, 2023.
[47] AirNow, https://www.airnow.gov/.