

1 **From consistency to flexibility: a database schema for the**

2 **management of CityJSON 3D City Models**

3 **Abstract:** The use of 3D city models is now common practice; many large cities have their
4 own digital model. Resilient and sustainable management of these models is necessary in
5 many cases, where an application could evolve over its life cycle. The complexity of generic
6 modelling standardization is often a limitation for a light and user-friendly usage and further
7 developments. This paper aims to propose an alternative providing a simplified database
8 schema implemented in a document-oriented storage. Thanks to the use of the NoSQL store,
9 the focus is on flexibility of the data schemas and thus its clarification. In order to aim
10 attention at the compactness in web development, CityJSON has been chosen for the
11 encoding of the 3D city models. Finally, a full-stack application (persistent storage,
12 consistent edition and visualization of 3D city models) has been developed to handle the
13 simplified schema and illustrates its capabilities in two practical use cases.

14 **Keywords:** CityJSON, NoSQL, 3D City Models, Data Schema, Data Architecture

15

16 **1. Introduction**

17 Nowadays, many large cities have usage of their own 3D digital model (Biljecki et al.
18 2015). These 3D city models are the integrating base for urban management tools such as
19 fluid flows simulations, cadastral operations, urbanism, etc. In the context of urban built
20 environment, the use of CityGML as the data model and encoding standard is now a common
21 practice (Gröger and Plümer 2012). CityGML provides a data exchange format for the
22 structuring of urban and landscape objects. It stores objects in multi levels-of-detail and
23 structures their attributes, their relationships and their features on a normalized basis. Its

24 support of an increasing number of extensions allows dealing with more and more issues:
25 energy, noise, land administration, etc. (Floros and Dimopoulou 2016; Biljecki, Kumar, and
26 Nagel 2018). From a conceptual viewpoint, these application domain extensions (ADE)
27 extend the supported features and properties of the CityGML core module. These added
28 elements are necessary to perform computations or to store their results in simulations and
29 analysis.

30 Recently, 3DCityDB, an open-source 3D geodatabase solution, has been proposed to
31 handle city models (Yao et al. 2018). The tool proposes a system for the management,
32 analysis, and visualization of large 3D city models according to the CityGML standard. It
33 relies on a relational database and provides well-known tools such as WFS services, the
34 support of 3D scenes (KML, COLLADA, etc), the streaming of these formats thanks to the
35 WFS capabilities, etc. The major drawback highlighted by the author states that the lack of
36 flexibility of the 3DCityDB relational solution could limit its usability; even if ADEs are
37 supported, maintaining them natively could be troublesome. Besides, the intrinsic
38 management of a relational solution might impose to make a large number of recursive joins
39 to represent the aggregation and inheritance hierarchies of the object-oriented data model.
40 Moreover, to support new features, it might be necessary to add tables, which always results
41 in an additional demand for resources and complexity of use.

42 This paper aims to provide an alternative to the relational database management of 3D
43 city model and traditional tools (SQL, CityGML, etc.). It relies on a simplified data schema
44 for the storage of city model in a document-oriented NoSQL store. A web three-tier
45 architecture (client, server and database), in which JavaScript articulates all the operations,
46 illustrates the use of the derived CityJSON schema, the JSON encoding of the CityGML data
47 model (Ledoux et al. 2019).

48 NoSQL databases offer the possibility to improve the storage flexibility by reforming
49 the tabular structure. Besides their reorganization of their intrinsic structure, this stores family

50 puts forward the plasticity of the schema model (Weglarz 2004). On the other hand,
51 CityJSON proposes a lightweight and compact alternative to the CityGML XML-encoding.
52 Following the same conceptual model as the XML-encoding, the JSON-encoding offers the
53 possibility to ease development of web applications. The conceptual similarities between
54 CityJSON and document-oriented management, which stores information as document in
55 BSON-encoding, could provide an answer to the lack of flexibility.

56 This paper is divided as follows: the section 2 contextualizes this research in related
57 works on Web Geographic Information Systems architecture (Web GIS) and the trend
58 towards an increasing use of the web (Mobasheri et al. 2020). It highlights the major
59 drawbacks of the current relational management and put it in parallel with the current state
60 of alternate developments. Then, the section 3 describes the simplified data schema and its
61 implementation in a document-oriented store. The illustrating application architecture is
62 decomposed in its three constituting parts: client, server and database. The section 4 develops
63 the new data management paradigm concerning the modifications provided by the NoSQL
64 database storage and several improvements on other tiers. A response is proposed and
65 documented in order to shed light on its new capabilities. From a network load viewpoint,
66 performances tests compare architecture capabilities in order to ensure exchanges
67 compactness. A benchmark with a relational solution is presented. Finally, two examples of
68 use cases illustrate these capabilities in practical situations in the section 5. Before
69 considering future works, we conclude on the principal benefits of the new generation
70 application and its advances.

71 **2. Related works**

72 A geographic information system (GIS) gathers and manages geospatial data (Tomlinson
73 1968). In the urban built environment, besides the management of 3D models and geometries,
74 the specific attributes and semantic information impose their own definitions; Urban GIS

75 (Blaschke et al. 2011). From a technical viewpoint, a web-based GIS application is divided
76 into three interdependent constituting parts at least: a client, which is a consumer of spatial
77 information; a server, which is a GIS processing system; and a database, which is a storage
78 solution that deals with spatial formats, spatial indexing and/or data processing functions. In
79 short, a Web GIS is a type of distributed information system in which components manage
80 spatial information on the web.

81 Nowadays, leveraging client capabilities and thus using its resources, the browser is no
82 longer simply a static window on a set of data: it can also perform a set of processes (Toschi et
83 al. 2017). Given that, the browser-based applications should outstrip standalone software
84 thanks to their multi-user characteristics and dynamic elements. It will result in cost savings
85 from the server without negative impact on the user experience (Kulawiak, Dawidowicz, and
86 Pacholczyk 2019). Indeed, the number of clients can also increase without limiting the server
87 performances, as it is used as a simple gateway and no longer as a computation centre.

88 Due to their mature support of spatial functions, indexes and storage capabilities, the
89 relational databases often represent the core base of web applications (Zlatanova and Stoter
90 2006; Mobasheri et al. 2020). Besides the data-modelling functions, the transactional
91 databases can handle data processing in an efficient way (Obe and Hsu 2015). Several
92 integrated solutions have been proposed for the management of digital city models. The
93 majority of these solutions are based on a relational database: (a) DB4GeO is a web service-
94 based geo-database architecture for geo-objects (Breunig et al. 2016). It relies on an object-
95 oriented database. Nevertheless, its development is no longer maintained. (b) 3DCityDB
96 provides a spatial relational database schema for semantic 3D city models (Yao et al. 2018).
97 It proposes an important number of key features and functionalities for CityGML models
98 management (Pispidikis and Dimopoulou 2016). It is interesting to note that, among other
99 functionalities, 3DCityDB allows the streaming of CityJSON features thanks to the OGC
100 WFS 2.0. (c) A NoSQL solution relies on a document-oriented storage and provides a 3D

101 web-rendering tool (Doboš and Steed 2012). However, these tools used in this architecture
102 were not as efficient as nowadays: many current libraries were unavailable (HTML5,
103 ThreeJS, etc.), the browsers capabilities were not as efficient as today; the focus was made
104 on the dataset and did not consider the architecture as a whole; etc. Moreover, the solution
105 developers criticized the lack of validation on elements import in the document-oriented
106 solutions. (d) Another NoSQL-solution development states that the document-oriented stores
107 lacks on consistency (Višnjevac et al. 2019). The problem here is that the database cannot itself
108 provide a sufficient guarantee of consistency. (e) The storage and manipulation of
109 heterogeneous data sources arises problems due to the differences in data structure: sensors
110 data, 3D city models, BIM models, etc. have a different update rate, a different representation
111 scale, etc. Even then, in GIS applications where sensors data, 3D city models and BIM
112 models coexist, the relational databases are preferred (Aleksandrov et al. 2019).

113 It is here worth mentioning that the dichotomy in which relational databases do not
114 support JSON insertion and document does is no longer true (Chasseur, Li, and Patel 2013).
115 Relational databases have been refactored to handle JSON (Liu, Hammerschmidt, and
116 McMahon 2014). However, it still imposes the use of an additional mapping layer and thus
117 does not provide a solution to the lack of flexibility. For instance, it is the case for 3DCityDB,
118 which translates the CityJSON in CityGML encoding before storing it into the relational
119 database thanks to the citygml4j software.

120 Developments on features visualisation have recently made progress on the client side
121 (Lim, Janssen, and Biljecki 2020). They provide a comparison on web-based viewers and
122 their specific capabilities at the building scale. However, the conclusions still draw the
123 disadvantages of ADE modelling and the complexity raised by relational database
124 management. Working on the storage tier, a composition of SQL/NoSQL allows enjoying
125 advantages of both solution (Holemans, Kasprzyk, and Donnay 2018; Poux et al. 2020).

126 While the relational database is still mandatory for its data-processing capabilities, the
127 document-oriented is useful thanks to its storage flexibility. It can be done without replication
128 or complex mapping between the two stores since the metadata and geo-registration are
129 handled on server side. The geospatial capabilities of the document-oriented stores bring
130 more and more solutions to spatial-related problematics (Zhang, Song, and Liu 2014; Lopez,
131 Couturier, and Lopez 2016; da Costa Rainho and Bernardino 2018). However, it shows that
132 even if performances are overall improved with document-oriented store, it is not yet always
133 true (Makris, Tserpes, and Anagnostopoulos 2019). Sometimes, relational database ranks
134 ahead of document-oriented stores (Bartoszewski, Piorkowski, and Lupa 2019), sometimes it
135 is the inverse in terms of loading (Laksono 2018) or heterogeneous sources handling (Sveen
136 2019).

137 From a technical viewpoint and in a more precisely way, MongoDB, a cross-platform
138 document-oriented database, has already been used in several “geo” architecture. Constituting
139 part of what is called a MERN stack (MongoDB - Express - React - NodeJS), MongoDB is
140 acknowledged for powerful way to store and retrieve data that allows developers to move
141 fast: MongoDB's horizontal, scale-out architecture can support huge volumes of both data
142 and traffic. Thanks to the flexibility of its database schema, this distribution has proved its
143 usefulness in spatial 2D (Đurić 2018; Voutos et al. 2017) and 3D visualization applications
144 (Trubka et al. 2016). The management of multiple representation structure can be visualized
145 using such a storage in the backend (Mao and Harrie 2016). However, its limited capabilities
146 to strict visualization could not set apart the document-oriented storages and its features.

147 About the stored data and the city modelling, CityJSON proposes to renew the CityGML
148 schema and provides a lightweight alternative to the XML encoding (Ledoux et al. 2019). Its
149 improved support of levels-of-detail and metadata make it a good substitute to CityGML
150 (Nys, Poux, and Billen 2020). However, its usage is still limited to specific applications and
151 data encoding (Kumar, Ledoux, and Stoter 2018; Nys, Billen, and Poux 2020; Virtanen et al.

2021). Besides it, the new support of 3D models in QGIS should improve its usability thanks to the development of a CityJSON plugin (Stelios Vitalis, Arroyo Otori, and Stoter 2020). Extensions of the core module are also promising way to improve the CityJSON usability and its update to the 3.0 CityGML version (Nys et al. 2021). In summary, nowadays, the storage of the CityJSON models are limited to files. There is currently no solution for storing and making models available in a collaborative and open manner.

3. Solution description

This section is divided in two subsections, respectively; a description of the simplified data schema for a document-oriented store and a description of the proposed architecture to demonstrate the usefulness of the proposed schema. While the first justify our choices on an efficient data accessibility and document nesting, the second is a short technical description of all the improvements made by an up-to-date WebGIS architecture.

3.1. Schema model

In the document-oriented database, the records are stored as documents that follow non-mandatory and semi-structured schemas (Olivera et al. 2015). All the documents respecting the same pre-established and semi-opened schema are gathered in a collection. These sets of documents allow the access and the indexing on the records or on a group of them. It is the primitive of the database query engine: everything revolves around this notion of collection. Note that, some efforts have been put to handle geospatial functions already but remain limited (Boaventura Filho et al. 2016). This section develops the various steps that led to enhance and modify the CityJSON encoding into a simplified database schema.

The bulk storage of a CityJSON city model in a single document without decomposing it in different collections is possible but limits the possibilities afterwards. A single collection storing all city models should therefore be queried as the document store works around this notion of set. Queries and indexing need to be complex to travel the embedded objects

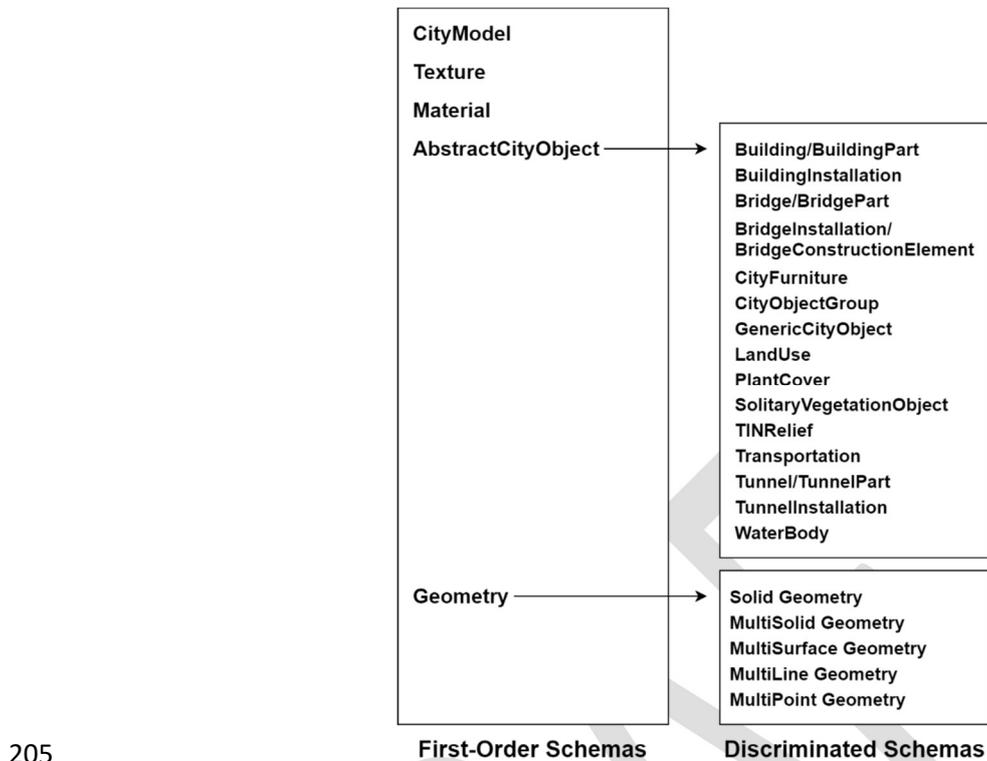
177 structure (an attribute is part of an object, which is itself part of the model). Even if compound
178 indexing is possible (i.e. successive levels of indexing on several attributes), this is not
179 recommended for efficient queries (Reis et al. 2018). Moreover, updating a sub-object in the
180 model without mobilizing the whole database become complex as it imposes to go deep in
181 the nondependent objects embedding, get the object and then insert the modified version in
182 the model.

183 Next to secondary elements such as metadata and appearances, a city model is made of
184 *CityObjects*. Those objects are natively embedded in the city model in a CityJSON file as
185 JSON objects. However, this data structure is not efficient enough for a dynamic use (Olivera
186 et al. 2015). According to the benchmark (Olivera et al. 2015), the referred models are more
187 efficient but impose to build dedicated queries. Consequently, once elements are created and
188 stored in collections, the link to referenced city objects need to be accessible from the city
189 model in a smart way.

190 We propose to create different collections in order to handle elements and ease their
191 access. Hence, we decompose the city model in five independent parts: *CityModel*, *Texture*,
192 *Material*, *AbstractCityObject* and *Geometry*. All imported records inherit their characteristics
193 from these five collections as their models are derived from these five top-schemas from the
194 CityJSON specifications (e.g. of a *Building* which is a specific *AbstractCityObject* with a n
195 *address*, a *measuredHeight*, a *roofType*, a specific set of allowed geometries, etc.). These
196 alternate schemas are the second-order schemas or discriminated schemas. In the core
197 application, the five first-order collections are defined dynamically by the database and the
198 server at startup (see Figure 1 **Error! Reference source not found.** for inheritance
199 relationships with second-order objects). Note that the *CityModel* collection represents the
200 models metadata only. A CityJSON model, as a file, is thus made of the gathering of its sub-
201 collections. Different models can be concurrently stored in the same database and the same
202 collections. Thanks to the database smart allocation of space, if a collection is empty, no

203 record is stored (i.e. collection does not exist at all, which implies that none space is used).

204 If a modification is made afterward, a new collection is created on the fly if necessary.



205

206

Figure 1. CityJSON objects schemas and inheritance

207 While importing the city model in the database, the city objects are stored as independent

208 objects in the *AbstractCityObjects* collection with a permanent link to their relative

209 *CityModel* document. Looping iteratively on the *CityObjects* array from the CityJSON file,

210 we create a new document for each new element and validate it depending on the city object

211 type (i.e. the validators are built on discriminated schemas independently according to the

212 CityJSON specifications and thus the CityGML data model). All elements are then stored in

213 the *CityObjects* collection whether it is a *Building*, one of its constituting *BuildingParts*, a

214 *SolitaryVegetationObject*, etc. In short, the schema imposes the necessary basis for files to

215 be correctly managed by the database and to follow the CityJSON core specification.

216 However, the management of this schema in a NoSQL solution does not limit the insertion

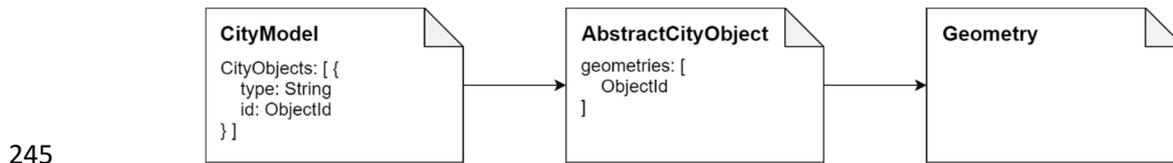
217 of extended attributes. Note that these extended attributes must still be coherent from a format

218 perspective: no special characters, no insertion functions, etc. Once a document is saved, its
219 corresponding document is afterwards referenced in the *CityModel* as a simplest object
220 stating on the type and the unique ID of the document in the *AbstractCityObject* collection
221 (see Figure 2**Error! Reference source not found.**).

222 As stated above, every object is referred with a unique identifier specific to its lifecycle
223 in the database (thanks to the special data type *ObjectID*). It is automatically generated and
224 indexed by the database. This integrated management allows concurrent users to create
225 objects at the same time but without any inconsistency insertion (i.e. users need to be aware
226 that two modifications can be made concurrently without any guarantee of consistency in a
227 NoSQL store). Note that the differences between the CityJSON discriminated schemas are
228 sometimes very subtle but this substructure allow further development in a convenient
229 manner: modification to the schema are easily made so that everything is decomposed,
230 normalized and structured. The addition of extensions takes direct advantage of this
231 flexibility as it might concern only a subschema or a part of it.

232 Concerning the insertion validation, during the model lifecycle, the *CityObjects* field can
233 therefore either be an entire object as in a file, either a reference or unique identifier to the
234 specific *CityObject* document. In order to prevent users to alter the consistency of the
235 database, it is thus important to provide a pivot element which can take one or the other value
236 without allowing too much deficiency (Diogo, Cabral, and Bernardino 2019). It imposes the
237 use of the *Mixed* datatype to validate the imported models. This pivot type is reused one more
238 time for the *CityObjects* to geometries relation (1-N relation). The Figure 2**Error! Reference**
239 **source not found.** illustrates the referenced structure of the first-order schemas in the
240 production phase; once documents have been created and referenced (i.e. value is fixed to
241 *ObjectID* and a string specifying the type of the object). In order to handle spatial indexing
242 and thus filtering queries responses spatially, a *geographicalExtent* attribute in computed

243 based on the geometry of every document. It corresponds to the smallest rectangular bounding
244 box enveloping the object geometries. This impacts performances on model import.



246 **Figure 2.** Referred documents structure in production

247 All geometries, and thus the fine and complex representation of the objects, are stored in
248 the same collection regardless of their type as has been the case with the city objects. Here,
249 it is not about a spatial management of elements (i.e. spatial functions and indexes are not
250 being used in the geometries collections) but about a management of elements of a spatial
251 nature (i.e. documents are actually real 3D objects following the standardized geometry
252 types). The geometries are complied with the ISO19107 standard according to the CityJSON
253 specifications. One more time, several discriminated schemas derive from the first-order
254 *Geometry* schema: *Solid*, *MultiSolid*, *MultiSurface*, *MultiLine* and *MultiPoint* (see Figure
255 **1Error! Reference source not found.**). Note that the “composite” geometries being
256 structurally similar to the “multi” ones, no new schema is created. They are managed as their
257 “multi” equivalent with the difference that their type is composite and not multiple. As a
258 reminder, the difference between the two is whether the constituent elements are contiguous
259 or not.

260 As in the CityJSON files (i.e. the Wavefront .obj file structure), the object boundaries
261 are stored as a list of vertices and arrays of pointers to vertices coordinate triplets in this list.
262 However, the referenced vertices triplets for every object are stored in bulk within the
263 *CityObject* document not in the whole *CityModel* one. This point set apart the database
264 schema with the common CityJSON files since the vertices should be stored in the *CityModel*
265 according to the specifications. In the direction of a wider support of spatial functions within

266 the application and the streaming of features, this storage method improves an independent
267 objects management: the spatial indexes and the consecutive references are suited for an
268 optimized spatial function support. Note that this discrete handling of vertices affect the
269 CityModel upload performances also. The support of spatial functions and tools represent an
270 important future work. Without tackling the database, it would also be interesting to consider
271 both server-side and client-side for spatial analysis.

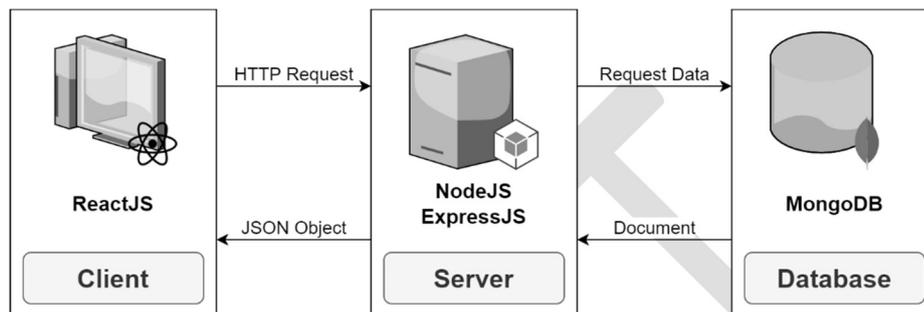
272 Concerning the support of schema extensions, an important benefit of the application
273 relates to the semi-openness of CityJSON specifications. While our motivation is to increase
274 flexibility, we would not limit the possibilities offered by the semi-open schemas. Hence, the
275 schemas structure is not locked. It allows the addition of attributes and/or properties and new
276 *CityObjects* type. We believe that CityJSON approach allow people to think about many
277 solutions in this way and ease their development. This point on total openness goes against
278 the 1.0.1 CityJSON specifications in which additional properties are not allowed in some
279 *CityObjects* definitions. Hence, some drawbacks might be encountered: an exported model
280 from the application might not be compliant with other tools in which specifications limit the
281 model to the strict conditions of the specifications. Efforts from the developers need to be
282 made in order to guarantee this interoperability.

283 3.2. *WebGIS architecture*

284 In the context of web development, when compactness and lightness are concerns, the
285 creation of a full-stack MERN (MongoDB - Express - React - NodeJS) app facilitates a smart
286 deployment. MERN web apps ensure convenience for web applications that have a large
287 amount of interactivity built into the front-end (i.e. the JavaScript clients). The following
288 paragraphs describe the constituting components of a MERN application and decomposes its
289 architecture in order to develop its benefits. Those benefits are mainly discussed concerning

290 their answer to the lack of flexibility of previous architecture and the availability of a database
291 support for CityJSON models.

292 Such kind of application is made up of a minimum of four technological stacks (ReactJS,
293 NodeJS, ExpressJS and MongoDB) as shown in Figure 3. The increase of flexibility and
294 resilience is demonstrated and put in parallel with the architecture components.



295

296 **Figure 3.** Architecture schema of a full stack MERN application

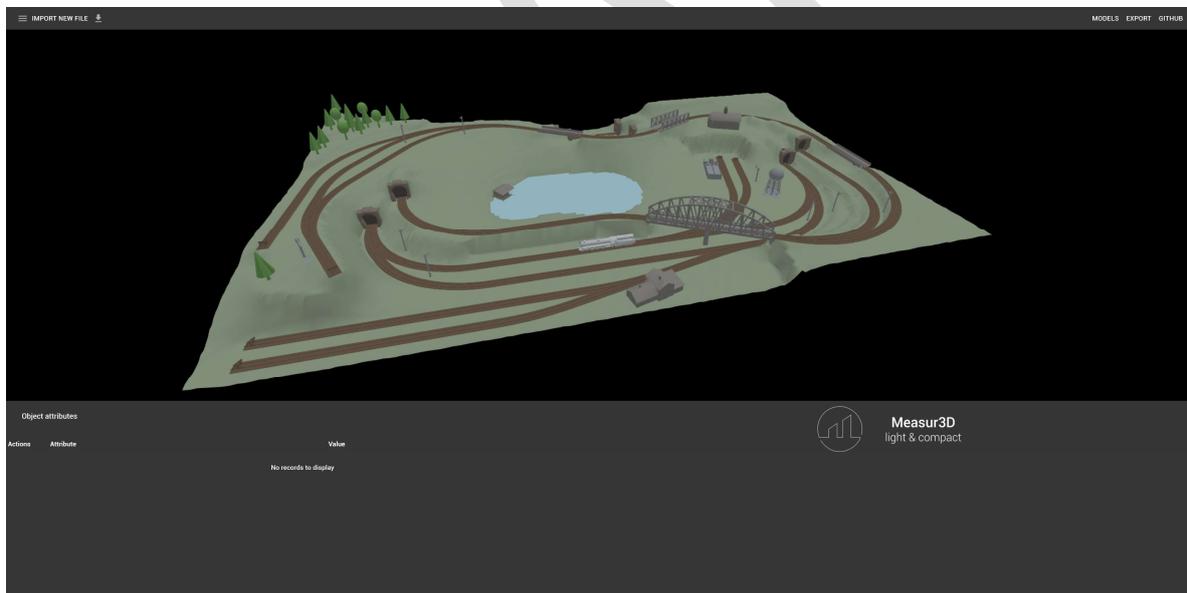
297 The four open-source constituting stacks of the core application are the following:

- 298 • MongoDB – the document-oriented NoSQL database.
- 299 • ExpressJS – a minimalist web framework for NodeJS.
- 300 • ReactJS – the Facebook MVC library (Model–View–Controller).
- 301 • NodeJS – a JavaScript runtime environment.

302 The client tier is built based on the ReactJS library (see Figure 4 for illustration). ReactJS
303 gave us the modularity necessary for the development of a new research tool as it does not
304 dictate a pattern. We thus focused on the data architecture and the application consistency. It
305 allows the construction of specific components and their reusability on a normalized basis.
306 Note that the rendering scene is an extension of the NINJA viewer (S. Vitalis et al. 2020). It
307 is itself based on the ThreeJS library (the WebGL cross-browser JavaScript library for 3D
308 manipulation and display). Nevertheless, the inserted value during updates and objects
309 modifications are tested in conformance with the *CityObject* schema and common insertion

310 rules (i.e. no special characters, no injections, etc.). The client tier allows all the common
311 CRUD operations (Create, Read, Update and Delete) on both *CityModels* and *CityObjects*.

312 The components communication is built on an event-driven paradigm: the components
313 subscribe to particular messages on an events bus. They then react to their subscription
314 whenever an update is published. The messages could carry information and/or simple
315 messages. It allows decoupling components in order to increase performance, reliability and
316 scalability (Allah Bukhsh, van Sinderen, and Singh 2015). Following this, all components
317 can be dismantled just as new components can be added modularly to open the application
318 possibilities. Hence, two panels are left open to integrate new modules for dedicated
319 functions: secondary view, tables, embedded objects, etc. Use cases of these panels are
320 presented in the end of this paper according to schema modifications during the production
321 phase.



322
323 **Figure 4.** Client view of the application – the rendered model is the dummy Railway.json
324 file provided by the 3D GeoInformation research group from TUDelft

325 The server is a NodeJS JavaScript runtime environment that allows performing
326 JavaScript code on server side (following the ECMAScript2015 specifications (Ecma

327 International 2015)). It follows an asynchronous, event-driven, non-blocking input/output
328 (I/O) model. These two last properties make it a very fast and resilient web server (Westerholt
329 and Resch 2015).

330 Along with that, ExpressJS is a JavaScript library that simplify the task of writing web
331 server code for NodeJS. Relying on HTTP requests (i.e. a RESTful application), it allows
332 server to set up middleware function calls: Cross-Origin Resource Sharing, rate limiter,
333 cache, compression, authentication, etc. Currently, the REST API performs basic functions
334 for CityJSON models and its features management such as CRUD functions. The
335 communication layer follows the HTTP/1.1 requests specifications. Please point out that the
336 non-successful responses are possible but non-response are avoided in conformity with the
337 BASE properties of the database. This property have been generalized to the server
338 application. Moreover, the server tier and thus the API ensure the application consistency as
339 the database itself does not provide any guarantee of it (Diogo et al. 2019).

340 The database tier is a document-oriented NoSQL store: MongoDB. Overall, the
341 document-oriented solutions tend to improve the performances and the storage volume for
342 dynamic data management. Despite many advantages, it is good remembering that the
343 responsibility to maintain the data sanity is no role of the NoSQL database (Diogo et al.
344 2019). The indexing method takes advantage of the metadata of each record. The choice of a
345 document-oriented solution has been made because of the schema flexibility and its native
346 JSON support (database object are BSON document of Binary-JSON object).

347 Unlike the English-like SQL, the dedicated MongoDB query language performs CRUD
348 functions but also aggregation, text search and a small number of geospatial queries. The
349 functions take JSON objects as parameters. Besides referenced relationships, the collections
350 are independent from one another. To make the comparison with relational databases, “joins”
351 are not allowed between collections. This point will be discussed in section 4.3.

352 4. Discussion on paradigm shift

353 Apart from the schema model and the proposed architecture, which have been discussed
354 on a technical aspect, several conceptual points need an explanation: the use of NoSQL was
355 not done without reason and some modifications to the CityGML/CityJSON conceptual
356 schema had to be made. The decomposition of the CityJSON files in documents and
357 collections schemas make up the structure of the database to perform normalized API calls.
358 This section comments the contribution of the simplified schema in order to open its reuse in
359 future works.

360 *4.1. Structured and unstructured data*

361 In this paper, we propose to shift the database archetype from relational solutions to a
362 NoSQL document-oriented store. This conversion should make it possible to open up
363 possibilities and ease schema modifications. While structured data (i.e. relational solutions)
364 promote a consistent data storage, unstructured data stores (i.e. NoSQL stores) intend to
365 enhance flexibility and availability.

366 The relational databases represent the more rigid storage structure. It imposes a static
367 tabular representation of the data (i.e. the data are imposed to follow a structure formatted as
368 rows and columns). The consistency of relational databases is especially ensured by the
369 respect of the ACID properties: Atomicity, Consistency, Isolation and Durability. The regard
370 of these properties results in the guarantee of avoiding insertion of inconsistencies in the
371 database. Conversely, the principal drawback of the relational family comes from the same
372 reason: the data querying and thus its availability can be slowed and inflexible because of all
373 the conditions imposed by ACID properties. Moreover, the table joins imposed by most
374 queries can make them cumbersome and result in complicated processes.

375 For instance, in the context of urban modelling, DB4GeO provides a solution relying on
376 an object-oriented database (OODB). Focusing on the data integrity, an OODB follows the

377 ACID properties. Even if the data structure established on objects is similar to NoSQL stores,
378 we find here the disadvantages of the relational model mentioned above. In addition, it is
379 difficult to make changes to an application that has been in production for some time. It
380 imposes to rework the database structure upstream, before any use. The section 5 illustrates
381 examples of how relational solutions need to be updated in order to handle new attributes
382 and/or new features using new associations.

383 Oppositely, in contrast with the rigid tabular models of relational databases, a document-
384 oriented store proposes to modify the data structure and open it. The NoSQL solutions do not
385 follow the ACID properties but the BASE properties (Basically Available, Soft state and
386 Eventual consistency). It results in a system in which denormalization is encouraged. The
387 horizontal scalability is improved (i.e. the replication of the system across n-database):

- 388 • Basically Available: the data are guaranteed as always available in terms of CAP
389 theorem. Whether it is successful or not, there is always a response to any request: “non-
390 response” are not possible from the store.
- 391 • Soft state: the state of the system could change over time. This can be possible even
392 without input. This is due of the eventually consistent property.
- 393 • Eventual consistency: the system will eventually become consistent once it stops
394 receiving input.

395 The document-oriented stores are composed of key-value pairs in which values can be
396 records such as XML, JSON objects or even other documents. For instance, sets of semi-
397 structured data might be deeply embedded and even recursive (i.e. chain references are
398 possible). Nevertheless, the management of records and lack of standardized schemas
399 improve their flexibility. It assumes a loss of records consistency to improve the database
400 flexibility because of the BASE properties. The consistency insurance is thus carried over to
401 server and client tiers and above all by the simplified schema. Here, the purpose is not on

402 the database consistency. A document-oriented store supports hierarchical documentation of
403 data, which is akin to CityJSON models and objects management. Every single records is
404 described by its own metadata. It uses agile and dynamic schemas without previously defined
405 structure.

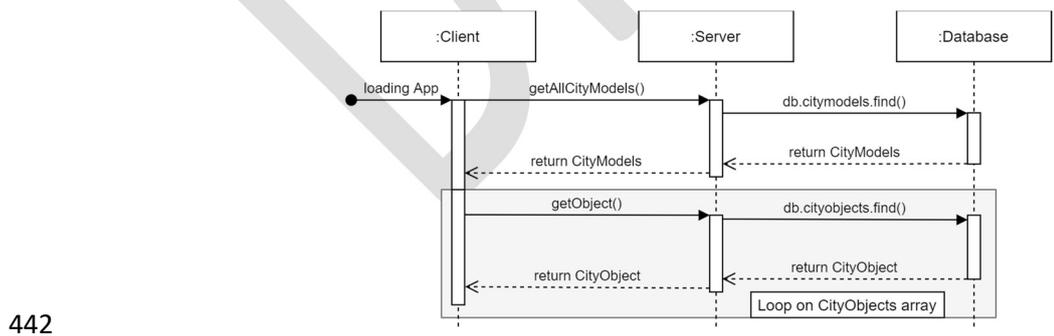
406 In summary, the alternative provided by the simplified database schema and its
407 implementation in document-oriented store allow users to ensure data availability and the
408 flexibility of their application in a simplified manner. It is not a solution that would go beyond
409 relational solutions but offers an opportunity to develop new functionalities. OGC API –
410 Features should indeed be an important improvement. It would take advantage of the
411 *CityObjects* collection, which corresponds to the notion of the standard: a set of features from
412 a dataset. Besides, the *CityObjects* are themselves abstractions of real world phenomena and
413 thus can be served as feature following the standard [ISO 19101-1:2014]. A discussion should
414 take place around these considerations and state on how CityJSON and the proposed
415 application can demonstrate it.

416 4.2. *Stacks communication*

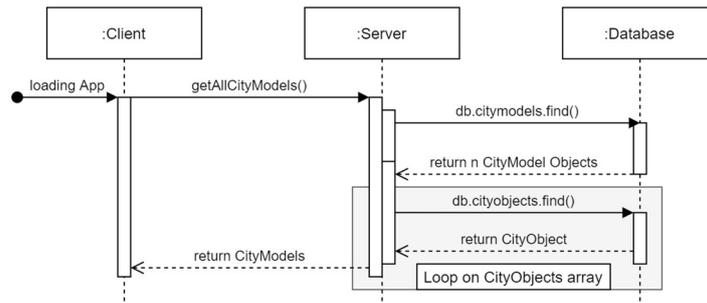
417 During the development of the application, while the client was hosted on a remote
418 machine, the application server and the database were hosted on the same machine. This
419 design allowed us to test server load, response time and response mode from a client/server
420 perspective. In order to assess on the best communication mode, we conducted tests on a city
421 model loading. The web GIS client capabilities becoming greater and greater (Agrawal and
422 Gupta 2017), we wanted to provide a benchmark of current objects managements possibilities
423 for a unique client (i.e. Chrome's V8 JavaScript engine in both server and client sides). Tests
424 in which n-clients queries the same API has also been made (see section 4.4). Downloading
425 the objects from the backend layer can be made in several ways:

- 426 • (a) Continuous requests: the server get all objects one by one from the database and send
 427 them to the client as soon as something is loaded. The city model reconstruction is carried
 428 by the client. It is characterised by a “flickering” apparition of elements in the rendering
 429 scene. It is a common asynchronous loading method.
- 430 • (b) Bulk requests: get all objects from the database then send them to the client in one
 431 aggregated object. The city model reconstruction is carried by the server. The model
 432 appears at once, in its entirety. It may take some time before seeing a result as all queries
 433 need to be resolved in order to response to the client.

434 Note that all exchanges are simplified thanks to the isomorphism of the application: all data
 435 are formatted as JSON objects in both back-end and front-end stacks. There is no need of
 436 translation or restructuring for the exchanges and the object management given that
 437 *CityObjects* are stored as they stand. In short, “what you store is what you access”. The Figure
 438 **5Error! Reference source not found.** and Figure **6Error! Reference source not found.**
 439 represent the sequence diagrams for both solution: continuous and bulk requests. They depict
 440 the succession of queries between the three-tier (client, server and database) and their
 441 responses.



443 **Figure 5.** (a) Continuous loading (sequence diagram) - client-side reconstruction.



444

445 **Figure 6. (b) Bulk loading (sequence diagram) - server-side reconstruction.**

446 The clients open a connection whenever they initialise themselves. The server and the
 447 database keep the connection open for future calls thanks to the NodeJS middleware. Hence,
 448 the client/server connection is made only once. Even if a client closes its connection, the
 449 database and the server keep a connection open for a limited amount of time in order to
 450 facilitate new connections. It is done given that opening a new connection takes a bit of time.

451 While the continuous loading allows diminishing the size of the bandwidth, the bulk
 452 loading allows making a single request on the network and reducing the global data transfer
 453 (i.e. fewer queries also means less redundancy in the formalization of query headers.).
 454 Moreover, caching the response of the bulk loading will improve performances as the model
 455 reconstruction is only made once. The tests were conducted on a small dataset, which
 456 numbers 120 *Building* objects and a *TINRelief* object. Note that, thanks to asynchrony from
 457 the NodeJS stack, the requests in the continuous loading were not stalled (i.e. no time were
 458 spent waiting because of proxy or ports negotiation before responses could be sent - the Time
 459 To First Byte (TTFB) was much nil). On the other hand, TTFB represented 99,6% of the bulk
 460 request time. It corresponds to the time for the server to process the database requests and
 461 reconstruct the whole city model before sending it. It is also important to note that time has
 462 been saved as *CityModels* are stored as they stand and thus the database does not need to
 463 formalize its responses. The whole process took twice as long for the continuous loading for
 464 a total amount of data exchanged four times greater (each request have a header and thus

465 multiply the size). Note that this consideration is only valid as long as the database structure
466 does not change.

467 4.3. *No joins*

468 Within a relational database, the objects are often split in several tables. Many
469 associations, which may be 1-1 but also 1-N and N-N cardinalities, link these tables together,
470 making it difficult to access the data. Modifying the stored objects, the number of relations
471 results in the modification of a potentially important number of tables. Moreover, this should
472 be done cascading in a specific order: first tables referred by foreign keys are modified, and
473 then tables linked with these specific keys. Hence, it is important to have a strong knowledge
474 of the database structure and provide guidelines and documentation to simplify developers
475 work.

476 On the other side, MongoDB retains the JSON objects structure and does not limit
477 insertions. For the reminder, this is not possible with a relational database that imposed the
478 use of conversion tools for native JSON file management. These tools often imply the
479 creation of many tables, many joins and thus the formalisation of complex queries. Such
480 queries and updates increase the time-consummation of processes due to the important
481 number of joins needed. Hence, if the conceptual model is complicated, it ends up with a lot
482 of complexity. A version attribute is modified on-the-fly allowing users to track elements.
483 The CityGML encoding is a perfect example of a high complexity structure (Yao et al. 2018).
484 For instance, in the 3DCityDB schema, sixty-six tables are used to handle CityGML models
485 in a relational database (against three collections in our simplified mapping and the use of
486 the *Mixed* datatype). The addition of modules increases this complexity but also might imply
487 to rework the database structure upstream. For instance, 3DCityDB and its import/export
488 tools allow creating new tables and associations in a convenient manner during the database
489 setup. Besides the addition of tables, it is worth specifying that these tables might be empty

490 or not use in practice: given that ADE are generic, all information might not exist or not be
491 relevant for the users' needs. This might be an additional source of bad resources
492 consumption. This is not the case in document-oriented solutions: empty fields simply does
493 not exist and documents structure evolves in accordance with the database lifecycle. In
494 summary, the repetitive joins, which are the main drawbacks of relational databases, are
495 avoided. This occurs in a more effective way to query, insert and store information whose
496 structure is assumed to change frequently. To compute results on several collections at the
497 same time, all collections need to be queried independently. The results are then gathered by
498 the client (e.g. of MapReduce processing techniques). As a reminder, the denormalization is
499 encouraged so reference and links can be done cleverly depending on the use of the product.

500 4.4. Comparison reference with relational solution

501 To illustrate the disadvantage of the relational joints, we conducted a benchmark on
502 several queries to 3DCityDB and our schema model. In order to perform these tests, we
503 simulated two remote JavaScript clients conducting queries on one side on a PostgreSQL
504 with the 3DCityDB model and on the other side on a MongoDB structured following our
505 schema. Both databases included the same three city datasets that counts 3471 objects in total
506 (3353 among them are *Buildings*). The query intends to get a random *Building* object with
507 its attributes (*roofType*, *function*, etc.), its unique ID and one of its *Solid* geometries.

508 Some elements need to be discussed before any statement. Before the instantiation, both
509 databases have a far different usage of memory. While 3DCityDB imposes the storage of 66
510 tables in 23Mb, our schema and its basic structure only takes 12Kb to create the three empty
511 collections. For the reminder, the collection schemas and the validation of an insertion are
512 handled by the server and not the database itself. It allows storage to be reduced and thus
513 improves performances. Once instanced, the relational solution is 149 Mb wide against 87Mb
514 for our schema (58%).

515 We have tested different interrogation methods by varying independently both the
 516 number of requests and the number of requested items. Note that the connection pool size of
 517 the database have an important impact on performances (a hundred was used). It is important
 518 to prepare it and to provide the same number of potential connections on both databases (by
 519 default, MongoDB allows only five concurrent connections. PostgreSQL allows hundred
 520 connections by default). It allows also to measure load under n-clients querying
 521 asynchronously the databases. About the architecture scalability, there is still room for
 522 improvement by multiplying the number of replicated databases (Schultz, Avitabile, and
 523 Cabral 2019). The balance should be determinate between the number of replications (n-
 524 databases), performance and the required consistency (Haughian, Osman, and Knottenbelt
 525 2016). Nevertheless, MongoDB offers already the possibility to create replications in a native
 526 way, which should facilitate future work.

527 As stated before, the relational schema imposes to inner join three tables. Our schema
 528 simply queries an object from the *CityObjects* collection specifying that the type of the
 529 queried object is "Building". Then it queries the related unique ID of the geometry in the
 530 *Geometries* collection. Since a document-oriented store is built and indexed on such relations
 531 and nested elements, this two steps retrieval seems to be more efficient. This hypothesis is
 532 directly reflected in the Table 1, which shows the databases response time.

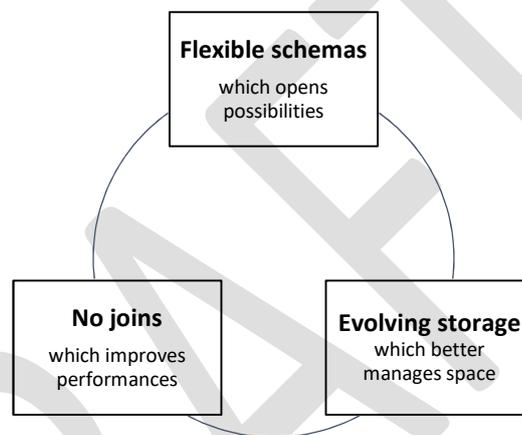
533 **Table 1.** Response time for the *Buildings* queries – repetition x objects (in milliseconds)

	1 x 1	1 x 10	10 x 1	1 x 100	100 x 1	1 x 3353 (1 x all)
Simplified schema	48	53	76	125	297	6678
3DCityDB	83	86	191	163	379	38089

534

535 These tests were conducted independently of the MERN application developments. In
 536 the application, a server cache avoids processing every query as some might be retained in
 537 the cache memory. In summary, this section offers an illustration of what is possible in the

538 matter of response time thanks to the new schema, the document-oriented storage and the
539 resilience of the MERN components. For the reminder, its contribution is a first answer to
540 the lack of flexibility of relational databases used in traditional architecture and the support
541 of CityJSON in a database. Hence, a convenient management of CityJSON models is thus
542 facilitated by the simplified schema, its three collections and the “what you store is what you
543 access” paradigm. A common base is given without limiting the usefulness of the schema to
544 a particular domain or specific end. These overall improvements of the schema and its
545 dedicated architecture can be summarized in three points (see Figure 7):



546

547

Figure 7. Summary of new capabilities

548 **5. Usage scenarios**

549 Now that the schema has been presented and the database solution has been compared
550 with a relational solution on a quantitative benchmark, we will state on the schema flexibility
551 through qualitative use cases. We have developed two simple extended schemas and two
552 modules to demonstrate the usefulness and the flexibility of the schema. It is illustrated in
553 situation of dynamic changes in the storage model during the production phase. The first one
554 is interested in the visualization of flat roofs and their potential for the installation of green
555 roofs. The second module concerns the management of the energy performance of buildings
556 certification and the updating of its calculation method. As a reminder, the structure of the

557 database is not modulated as the city objects are themselves not modified (collections are not
558 altered). However, the objects schemas allow the addition; the deletion and modification of
559 attributes in the stored records in a consistent way (see section 3.1).

560 *5.1. Urban green infrastructure*

561 Urban green infrastructures (UGIs) are part of the nature-based solutions for sustainable
562 urban development. In a previous research, we took part in the development of a simplistic
563 method for identifying the potential of green roofs along with identification of priority
564 regions in city centers (Joshi et al. 2020). In order to estimate the potential roof surfaces of
565 buildings, we interpolate planes based on a LiDAR point cloud and create building
566 geometries (Nys, Poux, et al. 2020). Once planes have been interpolated, we extract their
567 metrics such as the average heights of planes, their slope, their area, the number of planes per
568 buildings, etc.

569 During the method development, some limitations were noticed in a 2D framework
570 (Joshi et al. 2020): for instance, the obstructions are not considered (chimneys, elevator
571 shafts, etc). Taking into account a greater level of detail for the roof representation should
572 therefore improve the conclusion and catch the user's eye. As preparatory work for this new
573 study, we proposed to integrate the urban model into the application and add information as
574 it goes.

575 Therefore, we developed an extension that handles the relevant information for UGIs
576 installations. All information is attached to buildings geometries and integrated into the
577 CityJSON city model as object attributes. Besides, a modified version of the simplified
578 schema is hosted on the database. It validates the new attributes and guarantee the consistency
579 of the application through its different usages.

580 It was possible to add information relating to these levels of detail, whether purely
581 geometric or semantic, without modifying the work already done: the levels-of-detail

582 refinement were added to the model, even if it was already used by project partners. There
 583 was no need to create an additional collection. The visual report gives users a quick glance
 584 on the zone and future development solutions (see Figure 8). As stated in (Joshi et al. 2020),
 585 the method can still be improved considering more socio-economic factors. Hence, the
 586 application will allow handling the modifications easily and provides a convenient integrator
 587 basis for further developments.



588

589 **Figure 8.** UGI module for the visualization and computation of green roofs

590 For comparison purpose, the Table 2 has been updated to present response time of the
 591 *Building* query on the relational enhanced solution. In order to store the new information
 592 related to UGI, we added a table associated with the building one. Queries therefore impose
 593 the use of an additional join and thus affect performances, what we expected. Changes for
 594 the simplified queries in the NoSQL store are about the millisecond sometimes more,
 595 sometimes less. It has thus been not added to the table.

596 **Table 2.** Response time for the *Buildings* queries – repetition x objects (in milliseconds)

	1 x 1	1 x 10	10 x 1	1 x 100	100 x 1	1 x 3353 (1 x all)
Simplified schema	48	53	76	125	297	6678

3DCityDB	83	86	191	163	379	38089
3DCityDB + UGI	88	91	252	172	412	41374

597

598 *5.2. Energy performance of buildings*

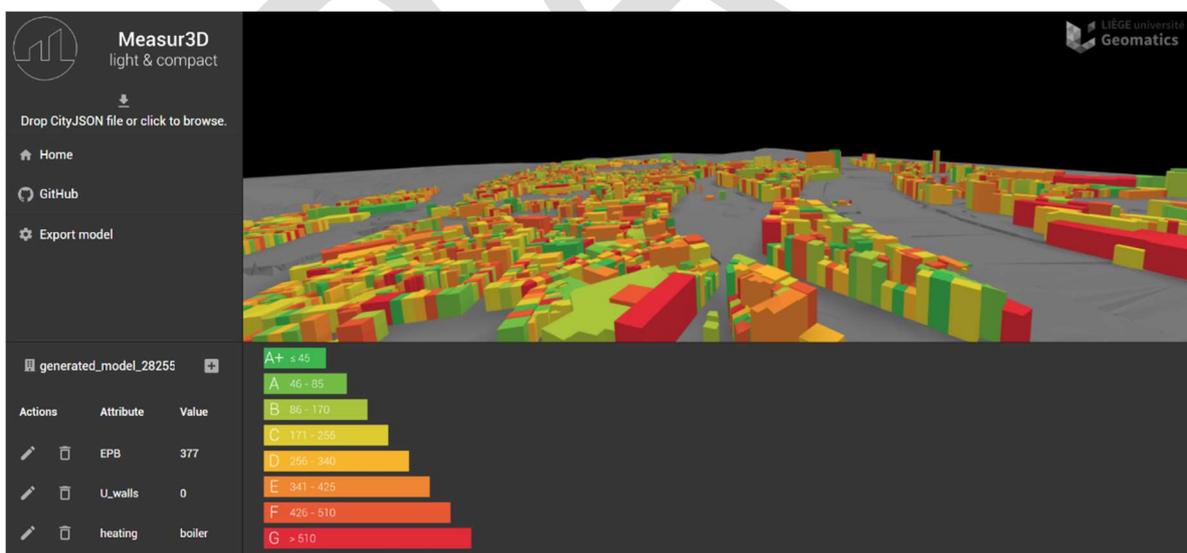
599 The European Directive 2010/31/EU of 19 May 2010 on the Energy Performance of
600 Buildings (EPB) requires Member States to set up a system of certification. In addition to
601 setting EPB requirements related to construction, it also imposes renovation work. The
602 energy performance certification of buildings consists of an overall assessment of the energy
603 performance of a building according to a defined calculation method.

604 In Belgium, this directive has been translated in an order of the regional government.
605 This order reviews the calculation method on occasion and makes changes at both the
606 semantic and conceptual levels. Depending on the modifications, the calculation of the
607 energy potential of buildings can change: new parameters can be included, some can be
608 deleted, new stats and intermediate values can be useful or neglected, etc. In an EPB
609 dedicated application based on a storage solution, all these statements result either in a
610 structure modification for new features either storing redundant, unnecessary or incomplete
611 information. As stated in the previous section, the usage of a NoSQL document-oriented
612 solution allows adapting the object attributes without any condition and storing them within
613 the same documents. This can be made without altering the database structure and frees
614 unused space as it goes.

615 The use of an architecture presented in this paper offers a flexible tool that can be easily
616 improved through different changes in methods and legislation. Without going into details of
617 the EPB calculation, we developed a module allowing calculating its value based on buildings
618 attributes and metrics. It is computed on the fly and changes buildings colour following the
619 normalised EPB scale (on the bottom left of Figure 9 - version updated on January 1, 2019).

620 The Figure 9 illustrates a simulation on 2369 buildings in the centre of Liège, Belgium. The
621 EPB module computes and stores the performance value based on attributes such as the type
622 of heating, the coefficient of thermal transmission of a wall, etc. We simulated a modification
623 in the EPB computation by taking into account the over-ventilation by manual opening of
624 doors and windows (in accordance with the decree of 11 April 2019). It was thus sufficient
625 to save the value but without modifying the database query mode using the REST API. The
626 database has thus added key/value pairs to the schema and the required documents in the
627 *Buildings* documents of the *AbstracCityObjets* collections.

628 The use of the tool proposes to handle both energy consumption data and 3D city models.
629 Rather than manage the certification on an individual basis, we offer the possibility to build
630 an energy cadastre at the neighbourhood scale but also of the city. The tool can be used by
631 communities for managing their energy consumption and perhaps optimizing them:
632 highlighting heat islands, heat plant installation, real estate renovation campaign, etc.



633

634 **Figure 9.** Illustration of the EPB module

635 6. Conclusion

636 This paper presents a simplified schema for the storage of 3D city model in a document-
637 oriented store. It illustrates new capabilities in a dedicated application that allows the storage,

638 management and visualization of CityJSON models. The JSON-encoding provided by the
639 CityJSON specifications has been opened and partially reworked in order to extend
640 possibilities of management. The different collections bring together the three main elements
641 of city models (*CityModel*, *CityObjects* and *Geometries*) and ensure data access. The
642 simplifications brought by this new model ease the accessibility and storage volume.

643 Besides, in order to demonstrate the capabilities of this simplified schema, we developed
644 an application based on JavaScript technological stacks and a NoSQL database. This database
645 paradigm shift proposes to go from a solution that ensure consistency (i.e. the ACID
646 properties of the relational databases) to a solution that improves the application flexibility
647 (i.e. the semi-openness of NoSQL schemas). The benchmark of this solution with the state
648 of the art is convincing in terms of response time and storage weight. We believe that this
649 application will improve the usage of CityJSON and web-based tools in urban built
650 environment modelling. The usability of the application has been illustrated in two use cases
651 of common practice: the visualization and the storage of urban green infrastructures and the
652 energy performance of buildings certification. The application allows users managing the
653 diverse data sources and structural changes during the production phase in a convenient
654 manner.

655 Future works will study the implementation of spatial functions support for the
656 application. An important discussion will take place on the choice between the three
657 possibilities of spatial support: database, client-side or server-side. While the former could
658 not be done without a deep rework of the database management, the proposed architecture
659 may have a place in the demonstration of spatial client/server capabilities enhancements.
660 Nevertheless, such improvements should keep an eye on the implementation of the OGC API
661 - Features standard in order to allow features fetching. A major improvement of this kind will
662 improve the user-friendliness and the dissemination of CityJSON models.

663 **References**

- 664 Agrawal, Sonam, and R. D. Gupta. 2017. 'Web GIS and Its Architecture: A Review'. *Arabian*
665 *Journal of Geosciences* 10(23):518. doi: 10.1007/s12517-017-3296-2.
- 666 Aleksandrov, M., A. Diakit , J. Yan, W. Li, and S. Zlatanova. 2019. 'SYSTEMS
667 ARCHITECTURE FOR MANAGEMENT OF BIM, 3D GIS AND SENSORS
668 DATA'. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information*
669 *Sciences IV-4/W9:3–10*. doi: 10.5194/isprs-annals-IV-4-W9-3-2019.
- 670 Allah Bukhsh, Zaharah, Marten van Sinderen, and P. M. Singh. 2015. 'SOA and EDA: A
671 Comparative Study - Similarities, Differences and Conceptual Guidelines on Their
672 Usage': Pp. 213–20 in *Proceedings of the 12th International Conference on e-*
673 *Business*. Colmar, Alsace, France: SCITEPRESS - Science and Technology
674 Publications.
- 675 Bartoszewski, Dominik, Adam Piorkowski, and Michal Lupa. 2019. 'The Comparison of
676 Processing Efficiency of Spatial Data for PostGIS and MongoDB Databases'. Pp.
677 291–302 in *Beyond Databases, Architectures and Structures. Paving the Road to*
678 *Smart Data Processing and Analysis*. Vol. 1018, *Communications in Computer and*
679 *Information Science*, edited by S. Kozielski, D. Mrozek, P. Kasprowski, B. Malysiak-
680 Mrozek, and D. Kostrzewa. Cham: Springer International Publishing.
- 681 Biljecki, Filip, Kavisha Kumar, and Claus Nagel. 2018. 'CityGML Application Domain
682 Extension (ADE): Overview of Developments'. *Open Geospatial Data, Software and*
683 *Standards* 3(1):13. doi: 10.1186/s40965-018-0055-6.
- 684 Biljecki, Filip, Jantien Stoter, Hugo Ledoux, Sisi Zlatanova, and Arzu  oltekin. 2015.
685 'Applications of 3D City Models: State of the Art Review'. *ISPRS International*
686 *Journal of Geo-Information* 4(4):2842–89. doi: 10.3390/ijgi4042842.
- 687 Blaschke, Thomas, Geoffrey J. Hay, Qihao Weng, and Bernd Resch. 2011. 'Collective
688 Sensing: Integrating Geospatial Technologies to Understand Urban Systems—An
689 Overview'. *Remote Sensing* 3(8):1743–76. doi: 10.3390/rs3081743.
- 690 Boaventura Filho, Wagner, Harley Vera Olivera, Maristela Holanda, and Aleteia Favacho.
691 2016. 'Geographic Data Modelling for NoSQL Document-Oriented Databases'.
692 Lisbon, Portugal.
- 693 Breunig, Martin, Paul V. Kuper, Edgar Butwilowski, Andreas Thomsen, Markus Jahn, Andr 
694 Dittrich, Mulhim Al-Doori, Darya Golovko, and Mathias Menninghaus. 2016. 'The
695 Story of DB4GeO – A Service-Based Geo-Database Architecture to Support Multi-
696 Dimensional Data Analysis and Visualization'. *ISPRS Journal of Photogrammetry*
697 *and Remote Sensing* 117:187–205. doi: 10.1016/j.isprsjprs.2015.12.006.

- 698 Brewer, Eric A. 2000. 'Towards Robust Distributed Systems'. P. 7 in *Proceedings of the*
699 *nineteenth annual ACM symposium on Principles of distributed computing - PODC*
700 *'00*. Portland, Oregon, United States: ACM Press.
- 701 Chasseur, Craig, Yinan Li, and Jignesh Patel. 2013. 'Enabling JSON Document Stores in
702 Relational Systems'. in *Proceedings of the 16th International Workshop on the Web*
703 *and Databases 2013*. New York, NY, USA.
- 704 da Costa Rainho, Filipe, and Jorge Bernardino. 2018. 'Web GIS: A New System to Store
705 Spatial Data Using GeoJSON in MongoDB'. Pp. 1–6 in *2018 13th Iberian*
706 *Conference on Information Systems and Technologies (CISTI)*. Caceres: IEEE.
- 707 Diogo, Miguel, Bruno Cabral, and Jorge Bernardino. 2019. 'Consistency Models of NoSQL
708 Databases'. *Future Internet* 11(2):43. doi: 10.3390/fi11020043.
- 709 Doboš, Jozef, and Anthony Steed. 2012. '3D Revision Control Framework'. P. 121 in
710 *Proceedings of the 17th International Conference on 3D Web Technology - Web3D*
711 *'12*. Los Angeles, California: ACM Press.
- 712 Đurić, Mladen. 2018. 'GEOPORTAL FOR SEARCHING AND VISUALIZATION OF
713 CADASTRAL DATA'. *CABPEMEHA TEOPIJA I IIPAKCA Y*
714 *ГPAДИТЕЉCTBY* 13(1). doi: 10.7251/STP1813687A.
- 715 Ecma International. 2015. *ECMAScript 2015. Language Specification*. 6. Geneva,
716 Switzerland.
- 717 Floros, G., and E. Dimopoulou. 2016. 'Investigating the Enrichment of a 3D City Model with
718 Various CityGML Modules'. *ISPRS - International Archives of the Photogrammetry,*
719 *Remote Sensing and Spatial Information Sciences XLII-2/W2:3–9*. doi:
720 10.5194/isprs-archives-XLII-2-W2-3-2016.
- 721 Gröger, Gerhard, and Lutz Plümer. 2012. 'CityGML – Interoperable Semantic 3D City
722 Models'. *ISPRS Journal of Photogrammetry and Remote Sensing* 71:12–33. doi:
723 10.1016/j.isprsjprs.2012.04.004.
- 724 Haughian, Gerard, Rasha Osman, and William J. Knottenbelt. 2016. 'Benchmarking
725 Replication in Cassandra and MongoDB NoSQL Datastores'. Pp. 152–66 in
726 *Database and Expert Systems Applications*. Vol. 9828, *Lecture Notes in Computer*
727 *Science*, edited by S. Hartmann and H. Ma. Cham: Springer International Publishing.
- 728 Holemans, Amandine, Jean-Paul Kasprzyk, and Jean-Paul Donnay. 2018. 'Coupling an
729 Unstructured NoSQL Database with a Geographic Information System'. Pp. 23–28
730 in *Proceedings of GEOProcessing 2018*. Rome, Italy.
- 731 Joshi, M. Y., W. Selmi, Marc Binard, Gilles-Antoine Nys, and Jacques Teller. 2020.
732 'POTENTIAL FOR URBAN GREENING WITH GREEN ROOFS: A WAY

- 733 TOWARDS SMART CITIES'. *ISPRS Annals of Photogrammetry, Remote Sensing*
734 *and Spatial Information Sciences* VI-4/W2-2020:87–94. doi: 10.5194/isprs-annals-
735 VI-4-W2-2020-87-2020.
- 736 Kulawiak, Marcin, Agnieszka Dawidowicz, and Marek Emanuel Pacholczyk. 2019.
737 'Analysis of Server-Side and Client-Side Web-GIS Data Processing Methods on the
738 Example of JTS and JSTS Using Open Data from OSM and Geoportal'. *Computers*
739 *& Geosciences* 129:26–37. doi: 10.1016/j.cageo.2019.04.011.
- 740 Kumar, K., H. Ledoux, and J. Stoter. 2018. 'Dynamic 3D Visualization of Floods: Case of
741 the Netherlands'. *ISPRS - International Archives of the Photogrammetry, Remote*
742 *Sensing and Spatial Information Sciences* XLII-4/W10:83–87. doi: 10.5194/isprs-
743 archives-XLII-4-W10-83-2018.
- 744 Laksono, Dany. 2018. 'Testing Spatial Data Deliverance in SQL and NoSQL Database Using
745 NodeJS Fullstack Web App'. Pp. 1–5 in *2018 4th International Conference on*
746 *Science and Technology (ICST)*. Yogyakarta: IEEE.
- 747 Ledoux, Hugo, Ken Arroyo Otori, Kavisha Kumar, Balázs Dukai, Anna Labetski, and
748 Stelios Vitalis. 2019. 'CityJSON: A Compact and Easy-to-Use Encoding of the
749 CityGML Data Model'. *ArXiv:1902.09155 [Cs]*.
- 750 Lim, J., P. Janssen, and F. Biljecki. 2020. 'VISUALISING DETAILED CITYGML AND
751 ADE AT THE BUILDING SCALE'. *ISPRS - International Archives of the*
752 *Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIV-4/W1-
753 2020:83–90. doi: 10.5194/isprs-archives-XLIV-4-W1-2020-83-2020.
- 754 Liu, Zhen Hua, Beda Hammerschmidt, and Doug McMahon. 2014. 'JSON Data
755 Management: Supporting Schema-Less Development in RDBMS'. Pp. 1247–58 in
756 *Proceedings of the 2014 ACM SIGMOD international conference on Management of*
757 *data - SIGMOD '14*. Snowbird, Utah, USA: ACM Press.
- 758 Lopez, M., S. Couturier, and J. Lopez. 2016. 'Integration of NoSQL Databases for Analyzing
759 Spatial Information in Geographic Information System'. Pp. 351–55 in *2016 8th*
760 *International Conference on Computational Intelligence and Communication*
761 *Networks (CICN)*. Tehri, India: IEEE.
- 762 Makris, Antonios, Konstantinos Tserpes, and Dimosthenis Anagnostopoulos. 2019.
763 'Performance Evaluation of MongoDB and PostgreSQL for Spatio-Temporal Data'.
764 P. 8 in *Proceedings of EDBT/ICDT 2019*. Lisbon, Portugal: CEUR-WS.org.
- 765 Mao, Bo, and Lars Harrie. 2016. 'Methodology for the Efficient Progressive Distribution and
766 Visualization of 3D Building Objects'. *ISPRS International Journal of Geo-*
767 *Information* 5(10):185. doi: 10.3390/ijgi5100185.

- 768 Mobasheri, Amin, Helena Mitasova, Markus Neteler, Alexander Singleton, Hugo Ledoux,
769 and Maria Antonia Brovelli. 2020. 'Highlighting Recent Trends in Open Source
770 Geospatial Science and Software'. *Transactions in GIS* 24(5):1141–46. doi:
771 10.1111/tgis.12703.
- 772 Nys, Gilles-Antoine, Roland Billen, and Florent Poux. 2020. 'AUTOMATIC 3D
773 BUILDINGS COMPACT RECONSTRUCTION FROM LIDAR POINT CLOUDS'.
774 *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial*
775 *Information Sciences* XLIII-B2-2020:473–78. doi: 10.5194/isprs-archives-XLIII-B2-
776 2020-473-2020.
- 777 Nys, Gilles-Antoine, Abderrazzaq Kharroubi, Florent Poux, and Roland Billen. 2021. 'AN
778 EXTENSION OF CITYJSON TO SUPPORT POINT CLOUDS'. in *Proceedings of*
779 *the XXIVth ISPRS Congress*. Nice, France.
- 780 Nys, Gilles-Antoine, Florent Poux, and Roland Billen. 2020. 'CityJSON Building Generation
781 from Airborne LiDAR 3D Point Clouds'. *ISPRS International Journal of Geo-*
782 *Information* 9(9):521. doi: 10.3390/ijgi9090521.
- 783 Obe, Regina O., and Leo S. Hsu. 2015. *PostGIS in Action*. Second edition. Shelter Island,
784 NY: Manning.
- 785 Olivera, Harley Vera, Maristela Holanda, Valeria Guimarães, Fernanda Hondo, and Wagner
786 Boaventura Filho. 2015. 'Data Modeling for NoSQL Document-Oriented Databases'.
787 Cusco, Peru.
- 788 Pispidikis, I., and E. Dimopoulou. 2016. 'DEVELOPMENT OF A 3D WEBGIS SYSTEM
789 FOR RETRIEVING AND VISUALIZING CITYGML DATA BASED ON THEIR
790 GEOMETRIC AND SEMANTIC CHARACTERISTICS BY USING FREE AND
791 OPEN SOURCE TECHNOLOGY'. *ISPRS Annals of the Photogrammetry, Remote*
792 *Sensing and Spatial Information Sciences* IV-2/W1:47–53. doi: 10.5194/isprs-annals-
793 IV-2-W1-47-2016.
- 794 Poux, Florent, Roland Billen, Jean-Paul Kasprzyk, Pierre-Henri Lefebvre, and Pierre Hallot.
795 2020. 'A Built Heritage Information System Based on Point Cloud Data: HIS-PC'.
796 *ISPRS International Journal of Geo-Information* 9(10):588. doi:
797 10.3390/ijgi9100588.
- 798 Reis, Debora G., Fabio S. Gasparoni, Maristela Holanda, Marcio Victorino, Marcelo Ladeira,
799 and Edward O. Ribeiro. 2018. 'An Evaluation of Data Model for NoSQL Document-
800 Based Databases'. Pp. 616–25 in *Trends and Advances in Information Systems and*
801 *Technologies*. Vol. 745, edited by Á. Rocha, H. Adeli, L. P. Reis, and S. Costanzo.
802 Cham: Springer International Publishing.

- 803 Schultz, William, Tess Avitabile, and Alyson Cabral. 2019. 'Tunable Consistency in
804 MongoDB'. *Proceedings of the VLDB Endowment* 12(12):2071–81. doi:
805 10.14778/3352063.3352125.
- 806 Sveen, Atle Frenvik. 2019. 'Efficient Storage of Heterogeneous Geospatial Data in Spatial
807 Databases'. *Journal of Big Data* 6(1):102. doi: 10.1186/s40537-019-0262-8.
- 808 Tomlinson, Roger. 1968. 'Geographic Information System for Regional Planning'. in *Papers*
809 *of a CSIRO Symposium*. GA Stewart.
- 810 Toschi, I., E. Nocerino, F. Remondino, A. Revolti, G. Soria, and S. Piffer. 2017.
811 'GEOSPATIAL DATA PROCESSING FOR 3D CITY MODEL GENERATION,
812 MANAGEMENT AND VISUALIZATION'. *ISPRS - International Archives of the*
813 *Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-1/W1:527–*
814 *34*. doi: 10.5194/isprs-archives-XLII-1-W1-527-2017.
- 815 Trubka, Roman, Stephen Glackin, Oliver Lade, and Chris Pettit. 2016. 'A Web-Based 3D
816 Visualisation and Assessment System for Urban Precinct Scenario Modelling'. *ISPRS*
817 *Journal of Photogrammetry and Remote Sensing* 117:175–86. doi:
818 10.1016/j.isprsjprs.2015.12.003.
- 819 Virtanen, Juho-Pekka, Kaisa Jaalama, Tuulia Puustinen, Arttu Julin, Juha Hyypä, and
820 Hannu Hyypä. 2021. 'Near Real-Time Semantic View Analysis of 3D City Models
821 in Web Browser'. *ISPRS International Journal of Geo-Information* 10(3):138. doi:
822 10.3390/ijgi10030138.
- 823 Višnjevac, Nenad, Rajica Mihajlović, Mladen Šoškić, Željko Cvijetinić, and Branislav
824 Bajat. 2019. 'Prototype of the 3D Cadastral System Based on a NoSQL Database and
825 a JavaScript Visualization Application'. *ISPRS International Journal of Geo-*
826 *Information* 8(5):227. doi: 10.3390/ijgi8050227.
- 827 Vitalis, S., A. Labetski, F. Boersma, F. Dahle, X. Li, K. Arroyo Otori, H. Ledoux, and J.
828 Stoter. 2020. 'CITYJSON + WEB = NINJA'. *ISPRS Annals of Photogrammetry,*
829 *Remote Sensing and Spatial Information Sciences VI-4/W1-2020:167–73*. doi:
830 10.5194/isprs-annals-VI-4-W1-2020-167-2020.
- 831 Vitalis, Stelios, Ken Arroyo Otori, and Jantien Stoter. 2020. 'CityJSON in QGIS:
832 Development of an Open-source Plugin'. *Transactions in GIS* tgis.12657. doi:
833 10.1111/tgis.12657.
- 834 Voutos, Yorghos, Phivos Mylonas, Evaggelos Spyrou, and Eleni Charou. 2017. 'A Social
835 Environmental Sensor Network Integrated within a Web GIS Platform'. *Journal of*
836 *Sensor and Actuator Networks* 6(4):27. doi: 10.3390/jsan6040027.
- 837 Weglarz, Geoffrey. 2004. 'Two Worlds of Data – Unstructured and Structured'. *DM Review*.

- 838 Westerholt, Rene, and Bernd Resch. 2015. 'Asynchronous Geospatial Processing: An Event-
839 Driven Push-Based Architecture for the OGC Web Processing Service: Push-Based
840 Async Geo-Processing with the OGC WPS'. *Transactions in GIS* 19(3):455–79. doi:
841 10.1111/tgis.12104.
- 842 Yao, Zhihang, Claus Nagel, Felix Kunde, György Hudra, Philipp Willkomm, Andreas
843 Donaubaauer, Thomas Adolphi, and Thomas H. Kolbe. 2018. '3DCityDB - a 3D
844 Geodatabase Solution for the Management, Analysis, and Visualization of Semantic
845 3D City Models Based on CityGML'. *Open Geospatial Data, Software and
846 Standards* 3(1). doi: 10.1186/s40965-018-0046-7.
- 847 Zhang, Xiaomin, Wei Song, and Liming Liu. 2014. 'An Implementation Approach to Store
848 GIS Spatial Data on NoSQL Database'. Pp. 1–5 in *2014 22nd International
849 Conference on Geoinformatics*. Kaohsiung, Taiwan: IEEE.
- 850 Zlatanova, Sisi, and Jantien Stoter. 2006. 'The Role of DBMS in the New Generation GIS
851 Architecture'. Pp. 155–80 in *Frontiers of Geographic Information Technology*, edited
852 by S. Rana and J. Sharma. Berlin/Heidelberg: Springer-Verlag.

853

854