

# Integration of Database Management Systems and Video Game Engines for Optimal and Sub-Optimal Pathfinding

Katrina Ziebarth  
Earlham College  
Richmond, Indiana, USA  
kszieba19@earlham.edu

## ABSTRACT

This study will examine the feasibility of more closely integrating databases with video games, with motivations being to take advantage of increased data locality, optimizations already implemented in database management systems, and preexisting solutions to a number of problems encountered by modern video games. Using O’Grady’s “Bringing Database Management Systems and Video Game Engines Together” as a base, this study will investigate how his approach can be both extended to the use of Weighted A\* for sub-optimal pathfinding and applied to pathfinding problems derived from commercial games.

## KEYWORDS

video games, database management systems, pathfinding, sub-optimal, SQL, OpenRA,

## 1 INTRODUCTION

The use of databases to improve the efficiency of pathfinding in video games has implications for the potential complexity of video games. It offers a way to reduce the cost of pathfinding, freeing resources for the other components of a video game. In addition, this approach may help to combat issues commonly faced by large multiplayer games.

Perhaps the most relevant work within this field is O’Grady’s 2021 dissertation. The need for this research can be seen in O’Grady’s observation that “it is not uncommon for business logic of any kind (including video games) to read all data from the database, compute a subset of the data or transformation thereof in an imperative programming language, and then write it back periodically.” As O’Grady notes, increasing data locality could potentially allow for much greater efficiency than is possible with that process. O’Grady also observes that “DBMSs have decades of development and optimization under their belt.” Additionally, O’Grady argues that modern video games, particularly large multiplayer ones that take place in persistent worlds, confront issues already addressed by database management systems, such as “what does it mean when two players want to manipulate the same part of the game world at the same time?”, “Can we efficiently find a subset of game elements that are relevant for a singular calculation, such as the objects that are currently visible to one player?”, and “When we need to do these operations over many game objects at once, how do we avoid swapping parts of the huge game world in and out of the main memory?” [14].

O’Grady attempts to demonstrate the feasibility of more closely incorporating relational database management systems by implementing standard components of game engines, including pathfinding, in SQL. His 2021 work includes implementations in SQL of A\*,

a booking system for A\*, and an iterative pathfinding method that O’Grady links to what he describes as “the set-based philosophy behind SQL, which excels at specifying operations on many elements at once.” O’Grady evaluates the performance of the first and third of these implementations using pathfinding which uses pgRouting, an extension for PostgreSQL, as a baseline. He specifically notes that his thesis “refrains from introducing a *Domain Specific Language (DSL)* to achieve the outlined tasks and uses only plain SQL without utilizing any imperative extensions” [14]. Gill explains a DSL as being “a special-purpose language, designed to encapsulate possible computations in a specific domain” [9].

This work departs from O’Grady’s in its examination of sub-optimal pathfinding. The reason for this is the lack of necessity of optimal paths in video games under most circumstances. As Botea et al. note, “In games, the optimality of solutions is not seen as a must-have feature. Suboptimal paths that look reasonable to a user are acceptable” [6]. Similarly, Gao et al. write “finding an optimal path is usually too costly for real-world applications, an alternate is to find some acceptable path by sacrificing some quality of the result. Such ideas are widely used in video games and serving robots” [8]. The admission of sub-optimal pathfinding algorithms allows for a wider range of approaches, potentially leading to increased efficiency in the production of reasonable solutions.

Another choice made by this work is to use various pathfinding problems derived from commercial games for evaluation of performance, where O’Grady used a single unidentified “map of a skirmish” [14], likely derived from Red Alert or another of OpenRA’s official mods [2]. In contrast, this approach offers the potential to provide further evidence for the feasibility of closer incorporation of database management systems in commercial game development.

## 2 RELATED WORK

### 2.1 Utilization of Databases to Improve Efficiency

Previous work exists in other fields concerning the utilization of databases to improve the efficiency of other systems. Bendre et al. explore how a spreadsheet can be unified with a relational database while preserving advantages of both, with their DataSpread having a spreadsheet front-end and a database back-end. In describing their tool, they write that “DataSpread retains all the advantages of spreadsheets, including ease of use, ad-hoc analysis and visualization capabilities, and a schema-free nature, while also adding the advantages of traditional relational databases, such as scalability and the ability to use arbitrary SQL to import, filter, or join external or internal tables and have the results appear in the spreadsheet” [5].

Similarly, in his dissertation, O’Grady explores the implementation in SQL of not only pathfinding, but also of AI and map generation [14]. O’Grady’s 2019 work also investigates the implementation of map generation in SQL. It “focuses on *real-time strategy* (RTS) games,” and like his 2021 dissertation, it uses the OpenRA engine [13].

## 2.2 Pathfinding in Video Games

A vast amount of relevant work exists in the field of pathfinding in video games. For the sake of brevity, this literature review will focus on subareas relevant to the objectives of this project.

**2.2.1 The Role of A\*.** Kapi identifies A\*, an optimal algorithm, as the most prominent algorithm in video game pathfinding [10]. Similarly, Abd Algfoor notes that A\* is among the most well-known search algorithms in games and robotics and that it “inspired many modified and improved algorithms” [4]. Rabin and Sturtevant even write that “All game developers understand that A\* is the pathfinding search algorithm of choice” [15]. Thus, it is unsurprising that O’Grady’s 2021 investigation of pathfinding focuses on A\* and variants based on it [14].

**2.2.2 Optimization.** Kapi offers an overview of various ways of optimizing pathfinding in video games, including choice of graph representation, modification of heuristic functions, and choice of data structure for implementation [10].

**2.2.3 Sub-Optimal Pathfinding.** In regards to sub-optimal pathfinding, Botea et al. identify as promising the idea of “combining compressed path databases with hierarchical abstraction” [6].

Returning to the subject of A\*, Rabin and Sturtevant, in their list of ways to optimize A\*, suggest using an inadmissible heuristic. It is their observation that “a small amount of overestimating has large benefits with very little noticeable nonoptimality” [15].

**2.2.4 Heuristics.** The three heuristics which O’Grady discusses for use with A\* on grid-based maps are Manhattan distance for Von Neumann neighborhoods and Euclidean distance and Chebyshev distance for Moore neighborhoods [14].

Botea et al. note that Manhattan distance and Octile distance are “simple, fast to compute, and reasonably accurate on many map topologies” [6].

**2.2.5 Environments.** O’Grady’s 2019 work uses the OpenRA engine [13], as does his 2021 dissertation [14]. In explaining the reasons for his choice of it for his dissertation, O’Grady gives as advantages that it is open source and under GNU General Public License, has a development community which can be contacted for inquiries, has all core components written in an object-oriented programming language, and has a “clear-cut set of mechanisms” for managing several actors due to its focus on real time strategy games [14].

Video games utilize a range of techniques for terrain discretization, with Nash and Koenig giving examples of games using regular grids, navigation meshes, and circle-based waypoint graphs [12].

Unfortunately, it was not possible to determine whether O’Grady’s experiments to measure performance in OpenRA in his 2021 dissertation used Von Neumann neighborhoods or Moore neighborhoods. O’Grady discusses both types of neighborhoods and uses in his code the boolean pseudo-function `neighboring(a, b)` [14]. It seems

probable that his goal in doing so was to make the code more general and avoid the need to present separate implementations for the two types of neighborhoods. However, O’Grady does state that “we assume maps to generally be decomposable into grids of arbitrary granularity” [14], confirming that these experiments dealt only with 2D maps.

**2.2.6 Existing Benchmarks.** A number of pathfinding benchmarks have been drawn from commercial games which utilize 2D grids [17]. For 3D voxel grids, a benchmark set from the game Warframe has been made available, as Brewer and Sturtevant explain. They note that “in comparison with 2D grid maps, relatively little work has been done on planning directly in a 3D space representation” [7]. For this reason, this project will focus on pathfinding on 2D grids.

## 2.3 Weighted A\*

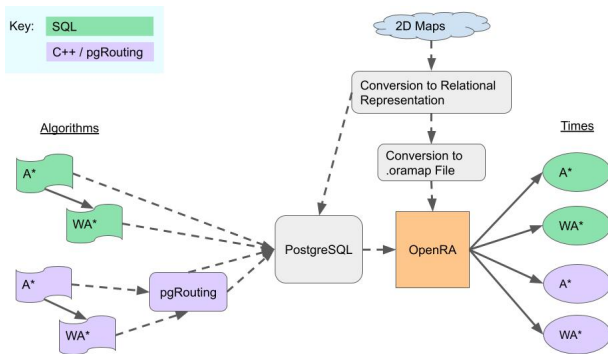
Rivera et al. note that Weighted A\* (WA\*) uses an evaluation function which is similar to A\*’s, except for its incorporation of a weight greater than or equal to one by which the heuristic function is multiplied. They also refer to weighting the heuristic as a “simple but powerful technique” [16]. This relates also to Rabin and Sturtevant’s suggestion that A\* be optimized for video games through use of an inadmissible heuristic. In fact, their approach of multiplying the heuristic portion of the formula for A\*’s evaluation by a weight is in practice identical to WA\* [15]. It is proposed by Rivera et al. that “A possible reason that explains why Weighted A\* finds solutions more quickly than regular A\* is that in multiplying the heuristic by a factor  $w \geq 1$ , the heuristic becomes more *accurate*, in a significant portion of the search space” [16].

## 2.4 pgRouting

As O’Grady notes, pgRouting is an extension which can be used with PostgreSQL “to offer path finding capabilities through a variety of path finding algorithms, including A\*, through calls to UDFs”. In his 2021 thesis, O’Grady uses pgRouting as a baseline for his work on pathfinding within a database management system (DBMS). While he did not find it compatible with “the ambition to fully realize components of video games in SQL in order to not be tied down to a specific DBMS,” due to its implementation in C++, he did mention that “it offers a way of keeping the computation of paths entirely in the world of the DBMS” [14]. pgRouting is also open source [1], which will allow for using a modified version of its implementation of A\* to implement Weighted A\* as well.

## 3 DESIGN AND IMPLEMENTATION

In accordance with the code given by O’Grady in his section on spatial A\*, with adjustments for omitted implementation details [14], A\* will be implemented in SQL. A modified version of the code will then be used in order to implement Weighted A\* in SQL. This decision was made due to the dominance of A\* in video game pathfinding, even in regards to optimization strategies. Similar to O’Grady’s 2021 work, pgRouting’s A\* implementation will also be used as a baseline, with Weighted A\* being implemented in pgRouting in order to allow for fairer comparison.



The performances of  $A^*$  and Weighted  $A^*$ , and pgRouting’s implementations of  $A^*$  and Weighted  $A^*$ , will then be compared on one or more sets of 2D pathfinding problems within OpenRA. These maps for these problems will be drawn from Sturtevant’s collection of existing benchmark sets derived from commercial games [17], with the specific ones chosen being derived from the real-time strategy games Warcraft III and StarCraft, due to the focus of OpenRA on real-time strategy. Since all of the maps within this collection are octile [17], neighborhoods within them would be Moore neighborhoods, and for this reason, Chebyshev distance, one of the two heuristics given by O’Grady for use with Moore neighborhoods [14], will be used as a common heuristic.

A Python program utilizing Psycopg 2 [3] will be used to convert the maps from a .map format to a relational representation in line with that described by O’Grady. Since the maps have already been processed to remove all but the largest connected component [17], it will not be necessary to check whether each map consists of a labyrinth in which paths exist between all random pairs of passable cells, as was true of the map which O’Grady used for evaluation [14]. Using OpenRA’s map editor, the maps will be converted to .oramap files, a format compatible with OpenRA, so that they can be used within the game engine. To obtain pathfinding problems for evaluation purposes, the converted maps will be used for skirmishes in which multiple AIs have been set to play each other. The same will be done for maps derived from Red Alert, which is already implemented in OpenRA.

In all experiments involving Weighted  $A^*$ , several weights will be used, in recognition of Rabin and Sturtevant’s observation that “The correct weight for your game or parts of your game must be discovered experimentally” [15], which implies that each set of pathfinding problems will require a different weight to optimize performance.

O’Grady used one action per 150 milliseconds as an loose upper bound for “agreeable waiting time” for pathfinding, citing Lewis et al.’s Starcraft study as a source [14]. Lewis et al. assert that professional Starcraft players in South Korea “can execute over 400 actions per minute (APM) in the game. By contrast, a highly accomplished amateur in the United States would likely top out in the mid 200s.” Lewis’ study, which utilizes replay files from games played in tournaments outside South Korea, uses a bound of 250 actions per minute [11]. This suggests that outside of high-performance contexts, 250 APM would likely yield a more reasonable upper bound

than the 400 APM O’Grady used. For that reason, this study will use metrics derived from both in evaluating algorithm performance.

All four algorithm implementations will be compared regarding time needed to calculate paths relative to path length as measured in nodes, similar to the comparison made by O’Grady between the pgRouting implementation of  $A^*$  and his implementation of  $A^*$  in SQL. In particular, emphasis will be placed on under what circumstances they remain within one or both of the upper bounds for waiting time. In recognition of O’Grady’s observation that “In real-world scenarios, especially for real time strategy games, short path searches are the most time-critical ones, as they are part of the micromanagement of units that occurs during heated battles” [14], heavier weight will be given to the results on shorter paths.

The relationship between optimality of path and path length in nodes will also be examined for Weighted  $A^*$ , since short path searches’ place as “part of the micromanagement of units that occurs during heated battles” [14] suggests they may be not only time-critical, but important to achieve near-optimality for.

## ACKNOWLEDGMENTS

I would like to thank David Barbella and Sofia Lemons for their assistance in planning this project. I would also like to thank Charlie Peck and Porter Libby for their assistance with this project. Additionally, I am grateful to Daniel O’Grady and Nathan R. Sturtevant for responding to my inquiries regarding their work.

## REFERENCES

- [1] 2020. *pgRouting Project*. Retrieved December 9, 2022 from <https://pgrouting.org/#>
- [2] 2023. *OpenRA*. Retrieved March 17, 2023 from <https://www.openra.net/>
- [3] 2023. *Psycopg 2*. Retrieved April 16, 2023 from <https://pypi.org/project/psycopg2/>
- [4] Zeyad Abd Algfoor, Mohd Shahrizal Sunar, and Hoshang Kolivand. 2015. A comprehensive study on pathfinding techniques for robotics and video games. *International Journal of Computer Games Technology* 2015 (2015), 1–11.
- [5] Mangesh Bendre, Bofan Sun, Ding Zhang, Xinyan Zhou, Kevin ChenChuan Chang, and Aditya Parameswaran. 2015. Dataspread: Unifying databases and spreadsheets. In *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases*, Vol. 8. NIH Public Access, 1–11.
- [6] Adi Botea, Bruno Bouzy, Michael Buro, Christian Bauckhage, and Dana Nau. 2013. Pathfinding in games. *Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik*, 21–31.
- [7] Daniel Brewer and Nathan R. Sturtevant. 2018. Benchmarks for Pathfinding in 3D Voxel Space. *Symposium on Combinatorial Search (SoCS)* (2018).
- [8] Zhipeng Gao, Junmeng Huang, and Chen Zhao. 2021. A double-phase search algorithm for sub-optimal path finding. In *2021 the 5th International Conference on Innovation in Artificial Intelligence*. 190–195.
- [9] Andy Gill. 2014. Domain-specific languages and code synthesis using Haskell. *Commun. ACM* 57, 6 (2014), 42–49.
- [10] Azyan Yusra Kapi, Mohd Shahrizal Sunar, and Muhamad Najib Zamri. 2020. A review on informed search algorithms for video games pathfinding. *International Journal* 9, 3 (2020).
- [11] Joshua Lewis, Patrick Trinh, and David Kirsh. 2011. A corpus analysis of strategy video game play in starcraft: Brood war. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 33.
- [12] Alex Nash and Sven Koenig. 2013. Any-angle path planning. *AI Magazine* 34, 4 (2013), 85–107.
- [13] Daniel O’Grady. 2019. Database-Supported Video Game Engines: Data-Driven Map Generation. *BTW 2019* (2019), 511–514.
- [14] Daniel O’Grady. 2021. *Bringing Database Management Systems and Video Game Engines Together*. Ph. D. Dissertation. Eberhard Karls Universität Tübingen.
- [15] Steve Rabin and Nathan R Sturtevant. 2013. Pathfinding architecture optimizations. *Game AI Pro: Collected Wisdom of Game AI Professionals* 1 (2013), 241–252.
- [16] Nicolás Rivera, Jorge A Baier, and Carlos Hernández. 2013. Weighted real-time heuristic search. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. 579–586.
- [17] Nathan R. Sturtevant. 2012. Benchmarks for Grid-Based Pathfinding. *Transactions on Computational Intelligence and AI in Games* 4, 2 (2012), 144 – 148. <http://web.cs.du.edu/~sturtevant/papers/benchmarks.pdf>