

Final Design Document - cc3k

Authors: Angelina Liang (a25liang), Derek Tan (d28tan), Paul Xiao (pxiao)

Overview

The game will first process the command line arguments, initialize input/output (I/O) resources, and enable the corresponding features based on the inputs.

Then, it would give the game engine `CC3K` access to those resources.

The game engine will create a menu for the player to select their desired race, then pass in the information to initialize a new game instance. Meanwhile, the instance of the menu is simultaneously destroyed.

The game will construct a new player based on the information given and initialize all in-game flags (e.g. `hostile_merchants`).

Then, it would start level generation per level, starting with the first level.

When the player goes up/downstairs, it will fetch the correct level from its stored list of levels to run. This design is to accommodate add-ons related to revisiting levels (floors).

The game will primarily run based on user input. It will pass the input command to the player, if the player makes a legitimate action, (i.e. the command is valid) then it will let enemies take their respective turns afterwards (unless the command for stopping movement has been issued).

The `level` class will contain all the information of that level (i.e. map, items, enemies) and is responsible for providing the related details (available locations, pointers to items/enemies) for the game to process.

The `map` class stores all map-related information (room numbers, walls, corridors, etc.) and handles random map generation.

If the player picks up any items (e.g. potions and treasures), the item ownership will be transferred over to the player class.

When a potion is applied to a character, the potion will be moved into the "effects" list of the character. The related stats of a character will get reset before each turn, and the corresponding potion effects will be applied to them to deduce the temporary stats for that turn.

This design is meant to allow players to apply potion effects onto enemies (via throwing potions) and to accommodate the turn-based duration of some additional DLC potions.

For attacks, the attacker will call the `get_hit()` method of the character being attacked, and both functions can be overridden by specific subclasses to implement custom special abilities.

Command Line Arguments

Optionally, the `cc3k` program can be run with additional command line arguments. You can run `./cc3k -h` or `./cc3k --help` to open the help information in the terminal.

Design

The following paragraphs describe the techniques used to solve the various design challenges encountered in this project.

For player character (PC) and enemy classes, we used the template method design pattern to create two abstract classes: `player_base` and `enemy_base`, and implemented each with a default (generic) set of interactions. To add the logic for a specific PC/enemy class, a subclass would inherit from the `player_base` or `enemy_base` class and override the logic to match the respective characteristics or abilities of each character.

Gold and potions both inherit from a common `item` class using the template method pattern. Potions are implemented with a decorator design pattern, such that their effects can be immediately applied by decorating the player class.

Abstract I/O classes also use the template method pattern to provide easy access to different I/O methods (`cin / cout`, file I/O, ncurses I/O).

At the start of each turn, the player stats are reset to accommodate the turn-based duration of potions. In addition, a timer is used to keep track of level-based duration, potions that last beyond the scope of a particular level.

For output, the output is not redisplayed upon every state change. Instead, it is written to a buffer and rendered all at once when the `out::render()` method is called.

If provided with a file for the five floors, when initializing the game (a CC3K object), it will preprocess all the given data to store information about the entities on the map and use a DFS algorithm to number the tiles for each room.

Random Map Generation is handled in the following manner:

- Generate the dimensions of all rooms
- Distribute the rooms horizontally into layers
- Distribute the layers vertically
- Jitter (randomly move upwards or downwards) the rooms vertically
- Generate the doors for each room
- Draw passages connecting the doors for all rooms

Ensuring that the player and stairwell do not spawn in the same room requires a way to identify the different rooms. Therefore, tiles in rooms are stored as numerical digits instead of '.' to keep track of the room number.

The ability to perform automated testing of the program is supported. Through a parsing algorithm, along with a `while` loop, the program can accept a text file of game commands as input, and then run the commands on a pre-specified seed to construct the same scenario for each run.

This way, despite the game having randomly generated components, it allows the reproducible, systematic testing of game logic and player interactions when new code changes are introduced.

Resilience to Change

There are some intentional design choices in the code structure that allow for the possibility of accommodating future changes to the program specification with minimal code modification and recompilation.

1) Layered Architecture

The project architecture delegates the task of presentation, business logic, and data access and retrieval amongst separate layers.

The presentation layer, consisting of the `input` class, `output` class, and their subclasses is responsible for processing and parsing all input data, as well as displaying the output in the desired format. This design allows the flexibility to add new output channels in the future (such as X11 graphics), as the output channel would only need to render the default output stream in a different style.

If the developer wishes to change the input syntax (for console-based input), the only file that would need to be changed is the `console_input.cc` file. Adding another input channel would only require creating another subclass of the `input` base class.

The presentation layer does not process any game logic, which instead is handled by the business logic layer, consisting of the `level`, `map`, and `character` base classes and corresponding subclasses.

Since all game interactions are handled by these classes in isolation, changes to the game logic (how different PC/enemy classes interact, and how potions affect others) would only require changes to these specific classes.

Data access and retrieval (e.g. in-game base stats) is owned and handled by the specific concrete subclasses of the player, enemy, and potion classes. All game-related stats are stored in the `constants.h` file, which is imported by the classes to extract the required data. This way, if a class' base stats were to be adjusted, any changes in the config file would automatically flow through without any additional refactoring.

This feature makes rebalancing characters and items much easier to perform. We used this functionality ourselves to rebalance some classes. The Monk class was nerfed due to initially being too powerful, and Mr. Goose was buffed to make it more survivable. Both changes solely required the tuning of their stats defined in the `constants.h` file.

This modular design promotes low coupling of modules since the presentation, business logic processing, and data access functionality are all separated into distinct components. Furthermore, all data related to a class is owned and handled by the class itself, preventing any co-mingling of data which can lead to unintentional data corruption.

This design also promotes high cohesion, since all modules related to a specific game function (e.g. player, enemy, potion) will inherit from the respective base class. Thus, all player character related classes inherit from `player_base`, all enemies inherit from `enemy_base`, etc.

There are a few utility libraries that are used across multiple different classes. These include: `rng`, `position`, and `fraction`. These tools have been abstracted away from the client classes to reduce the coupling of their common functionality and improve cohesion among all game-related classes. They are treated equivalently to external libraries in the code.

2) Interfaces, Abstraction, and Design Patterns

The interfaces for all modules are abstracted away from their implementation files (i.e. the interface files do not provide any in-line implementation). This is done so that the underlying implementation can be swapped out in the future without damaging the interface.

If the developer wishes to add a new player character, enemy, potion, or I/O channel into the game, they can override the virtual functions in the template class to have easy access to the individual features at an abstract level. This is done using the virtual constructor pattern for initializing the subclass, and then using the template method pattern to override the default functions of the base class. This approach applies to all player characters, enemies, potion, and I/O classes.

In addition, if the developer wishes to change the rules in the game, they can modify the game logic directly in `level.cc`. If they wish to change how an interaction occurs between two specific character classes (e.g. elf vs leprechaun), they only have to modify the implementations of the two subclasses (in this case, class `elf` and `leprechaun`). The same approach is used to modify player/potion, enemy/potion, and potion/potion interactions.

3) Global Configuration

Configurable aspects of the program are not stored within classes themselves, but rather in an external config file.

The unified constants file `constants.h` tracks the state of all global constants used in the cc3k program. These include player stats, enemy stats, potion stats, feature flags, and map stats. The developer would only need to edit this file to make any universal changes to the stats used for the game. This is much more effective than having such constants nested within classes, changing all of them manually would prove to be a logistical nightmare.

Furthermore, the program supports the ability to extend the current offering of command-line arguments. Any new command line arguments or opt-in features can be turned on/off by adding a feature flag to `arguments.cc` and updating the feature flags section in `constants.h`. Currently, additional command line arguments are available in the current program, run `./cc3k -h` to view its documentation.

Answers to Questions

Question 1: How could you design your system so that each race could be easily generated? Additionally, how difficult does such a solution make adding additional races?

Answer: In our project, two template classes handle races: `player_base` and `enemy_base`. The abstract character and enemy classes will contain common attributes among all classes, such as attack (ATK), defense, (DEF), health points (HP), and hit rates. The template classes also include implementations for the default basic interactions, such as attacking and getting hit.

Player character (PC) races (e.g. shade, drow, vampire, troll, goblin) will be subclasses of the class `player_base`. Enemy character races (e.g. human, dwarf, elf, orc, merchant, dragon, halfling) will be subclasses of class `enemy_base`. By design, it is manageable to add additional DLC classes to the game without refactoring the logic of the player character and enemy classes of the base game.

Question 2: How does your system handle generating different enemies? Is it different from how you generate the player character? Why or why not?

Answer: The system handles initializing different enemies to player characters in a similar manner. For enemy classes, the new enemy instance is inherited from the class `enemy_base` and the virtual functions of `enemy_base` are overridden in the implementation of each concrete enemy class.

The difference in generating between the player character and enemy classes is that enemy classes rely on a random number generator (RNG) object with predefined odds specified in the document, whereas the player character class generates with 100% certainty the selected player character class.

The enemy classes are designed this way to make it easy to add new enemy classes in the future (as part of a DLC), all the developer would need to do is to inherit from `enemy_base` and add the custom functionality in its implementation file.

Question 3: How could you implement the various abilities of the enemy characters? Do you use the same techniques as for the player character races? Explain.

Answer: Enemies have a set of basic abilities (normal attacks, normal movement, getting hit by others) in their template base class `enemy_base`, all of which are virtual and can be overridden.

Each basic ability in the `enemy_base` class contains a default implementation, such that if a new enemy class does not change its default functionality, then the virtual function can simply be inherited and not overridden (i.e. non-pure virtual function).

The same technique will be applied to the abilities of PC races (attack, applying potions, getting hit, etc.). For example, the basic set of operations for PCs is based on the shade race (i.e. shade doesn't need to override any method of `player_base` and is a generic PC class) and the vampire will override the attack function to include HP draining (i.e. lifesteal).

The `player_base` class also contains a set of default ability implementations that may be overridden or simply inherited as is, depending on the specific player character's race.

Question 4: What design pattern could you use to model the effects of temporary potions (Wound/Boost Atk/Def) so that you do not need to explicitly track which potions the player character has consumed on any particular floor?

Answer: We would use the Decorator design pattern to temporarily attach the effects of temporary potions (WA/WD/BA/BD) to the PC class. Immediately upon potion use, the decorators will directly alter the ATK/DEF/HP attributes of the player, thus the potion does not need to be explicitly tracked for its effects to be reflected in the player's stats.

At the end of each level, there will be a reset function that will reset the effects of all temporary potions used on the current floor, thus clearing its effects before advancing to the next floor. After the level is complete, all player character stats will reset to their default values (except HP), and a new round of interpreting will begin at the new level.

Question 5: How could you generate items so that the generation of Treasure and Potions reuses as much code as possible? That is, how would you structure your system so that the generation of a potion and then the generation of treasure does not duplicate code?

Answer: Both potions and gold (treasure) will inherit from a base class `item`. The generation function will have shared functionality in generating positions, and different type indicators can be inserted to create items of various flavours (i.e. different potion types and different gold types).

Since potions and gold inherit from the base class `item`, any possible duplication of code is avoided by the sharing of common functionality (e.g. position generation). This makes it a more efficient process when adding new potion/gold types, as well as improving the developer experience when maintaining different items in the game (since core item functionality is centralized in a single location).

Extra Credit Features

No Raw Pointers

What we did: All memory in the program is allocated using STL containers and smart pointers, and no raw pointers are used to manually allocate memory.

`ncurses` I/O

What we did: Print output using the `ncurses` library, which reduces latency.

Random Map Generation

What we did: Randomly generates a new map each time.

Why it was challenging: Random generation can get out of hand since there is no way to predict the generated output.

How we solved it: Setting bounds on the parameters used for generation (max/min room height/width, room spacing, reserved space).

Enemy Chasing

What we did: Enemies move towards the player when within a certain radius.

Transient Enemies

What we did: Enemies can walk through doors/passages.

Inventory

What we did: The player can store potions and items in their inventory (across levels).

Why it was challenging: Managing potion ownership.

How we solved it: Levels can release the ownership of individual potions and transfer it to the player class.

Throwable Potions

What we did: The player can throw a potion in their inventory for up to a certain distance.

Additional Player Characters

What we did: Added new playable character classes to the game, each with their unique characteristics and abilities.

The descriptions and special abilities of additional player characters are displayed by running the `./cc3k --races` command. (Try it!)

How we solved it: Override individual members of the `player_base` class.

Additional Enemies

What we did: Added new enemy classes to the game, each with their unique characteristics and abilities.

The descriptions and special abilities of additional enemy classes are displayed by running the `./cc3k --enemies` command. (Try it!)

How we did it: Override individual members of the `enemy_base` class.

Additional Potions

What we did: Added additional potion classes that had a variety of different effects not included in the base game.

The descriptions and special abilities of additional potions are displayed by running the `./cc3k --potions` command. (Try it!)

Why it was challenging: Additional potions have turn-based duration instead of lasting for an entire level.

How we solved it: All potions have an internal duration timer and we explicitly stored potions in the ownership of the character owning the effect.

Final Questions

Q1: What lessons did this project teach you about developing software in teams? If you worked alone, what lessons did you learn about writing large programs?

A: Many important lessons were learned throughout the course of doing this project. Specifically, it reminded us of the importance of version control, project/task tracking, and communication.

Importance of Version Control:

When simultaneously collaborating on a project with multiple contributors, we discovered that it is important to keep a single, centralized source of truth available at all times. For our group, using Git allowed us to have a stable working version of the code in the `master` branch at all times, as well as roll back any breaking changes to a previous working version of the code.

It helped us keep track of the latest changes to the code using commit messages and provided us with the ability to asynchronously work on different features without interfering with the master branch (through Git branching) or with each other.

Importance of Project and Task Tracking:

We also realized the importance of having a system to help us keep track of the current status of the project and the status of all pending tasks. This required a project management tool such as Trello to monitor and track the progress of such tasks.

Through a Kanban Board on Trello, we tracked the status of all tasks, bugs, chores, and features by bucketing them into the categories of To-do, Doing, Blocked, Under Review, and Done. Furthermore, we could assign tasks to individual members to clarify the work distribution.

Importance of Communication:

Through this process, we realized that communication was an essential pillar of effective group collaboration. In our group, we maintained a searchable chat log that we would use to communicate, and this helped us organize our thoughts and ideas. Whenever the code or design was updated, a message was sent to the group notifying everyone of the changes in detail. This practice helped our group avoid potential miscommunications, as well as build a traceable knowledge base for future reference. We also constantly updated the Trello board and informed group members of any blockers/updates when they occurred. Through these processes, we were able to constantly stay on top of tasks and always maintained a clear understanding of what needed to be done.

Q2: What would you have done differently if you had the chance to start over?

A:

1. **More time spent on planning and design:** Many of our features were inspired while working on the project, and we overlooked a few requirements at the beginning.
2. **Unify coding style:** While working, we found that individual members have different styles for indentations, naming, code structure, etc. Even though we managed to present a unified-looking codebase with the help of tools like AStyle, it required a lot of additional communication to help understand each other's code.

3. **Continuous testing:** We started out implementing individual modules and only writing testing harnesses specific to the modules. When we pieced them together after individual testing, we found that they did not coherently fit together. Moving forward, we should test both individual components (unit testing) as well as their interactions with other components (integration testing).
4. **Emphasis on documentation:** We had several loose files describing how each module and the entirety of the code worked but lacked any general documentation to explain the code in its entirety. We should document what we were doing with each piece of code along with how we did it.
5. **More robust structure:** We designed the code with some of the extra features in mind, but the features we implemented after required us to implement additional interfaces to accommodate. In the future, we should design interfaces with accessibility and modularity in mind.

Conclusion

As a concluding note, we would like to express our heartfelt gratitude to the CS246 instructors Caroline Kierstead and Jens Schmitz, as well as the CS246 instructional support staff, for their invaluable guidance, insightful advice, and unwavering support throughout this process.

Many DLC features in this project were inspired by NetHack, an open-source, single-player rogue-like video game also inspired by the 1980 video game Rogue.

This project has been a rewarding and enjoyable experience, through which we gained substantial knowledge in C++ and object-oriented programming. The skills and understanding we acquired in this project will undoubtedly serve us well in our future programming endeavours.

After careful consideration, we have decided not to become Shades¹ and instead transform our lives into those of Hackers². This project has received great help from Mr. Goose³ during debugging, which is a good alternative to trolling⁴.

This message was sponsored by your local neighbourhood leprechaun⁵.

Sincerely,

Angelina, Derek, and Paul

¹ "You either die a hero or live long enough to see yourself become the villain" - Harvey Dent, in the film The Dark Knight

² Hackers will attack Shades in the cc3k game.

³ The one who knows all.

⁴ Troll was the most broken character in the basic game, thus Mr. Goose is a good alternative.

⁵ They probably stole all your gold when you weren't looking.