Automated Planning to Prioritise Digital Forensics Investigation Cases Containing Indecent Images of Children

Saad Khan[1], Simon Parkinson[1], Monika Roopak[1], Rachel Armitage[2], Andrew Barlow[3], Ian Taylor[3], Matt Sawers[3], Dan Reiter[3]

[1] Department of Computer Science, [2] School of Human and Health Sciences, University of Huddersfield, Queensgate, Huddersfield HD1 3DH UK

[3] Kursch Consult Ltd, UK

saad.khan@hud.ac.uk

Abstract

Law enforcement agencies (LEAs) globally are facing high demand to view, process, and analyse digital evidence. Arrests for Indecent Images of Children (IIOC) have risen by a factor of 25 over the previous decade. A case typically requires the use of computing resources for between 2-4 weeks. The lengthy time is due to the sequential ordering of acquiring a forensically sound copy of all data, systematically extracting all images, before finally analysing each to automatically identify instances of known IIOC images (second-generation) or manually identifying new images (first-generation). It is therefore normal practice that an understanding of the image content is only obtained right at the end of the investigative process. A reduction in processing time would have a transformative impact, by enabling timely identification of victims, swift intervention with perpetrators to prevent re-offending, and reducing the traumatic psychological effects of any ongoing investigation for the accused and their families.

In this paper, a new approach to the digital forensic processes containing suspected IIOC content is presented, whereby in-process metrics are used to prioritise case handling, ensuring cases with a high probability of containing IIOC content are prioritised. The use of automated planning (AP) enables a systematic approach to case priorisation. In this paper, a planning approach is presented where AP is used to generate investigative actions in 60-minute segments, before re-planning to account for discoveries made during the execution of planned actions. A case study is provided consisting of 5 benchmark cases, demonstrating on average a reduction of 36% in processing time and a 26% reduction in time required to discover IIOC content.

Introduction

Both in the UK and globally, Police forces and law enforcement agencies (LEAs) are facing huge demand to view, process, and analyse digital evidence. The continuous increase in computer, mobile, and internet device use has created rapid growth in digital policing requirements. Future, the increase in case size and diversity in technology utilised by perpetrators is providing problematic for law enforcement to handle (Ayers 2009; Garfinkel 2010). In the United Kingdom, just as globally, the number of law enforcement cases suspected of containing Indecent Images of Children (IIOC) is increasing. Details on the specific number of IIOC cases processed in the UK are not available in isolation; however, the Office for National Statistics provides crime trend information in England and Wales over a 20-year duration. In the data set, the number of offences categorised as "Obscene publications, etc and protected sexual material"[1] is provided. Figure 1 presents the number of police-recorded offences involving indecent images of children, per year since 2012. As evident, the cases have rapidly increased by over a factor of six. This evidences an increased burden on law enforcement resources. Any delay in discovering IIOC content is preventing timely conviction and the opportunity to identify victims, which could permit ensuring their safety and well-being. Previous work including partners and children of offenders reveals that families can be without electronic devices for up to three years, adding to the psychological, social, and physical impact of the offense (Winder et al. 2020).

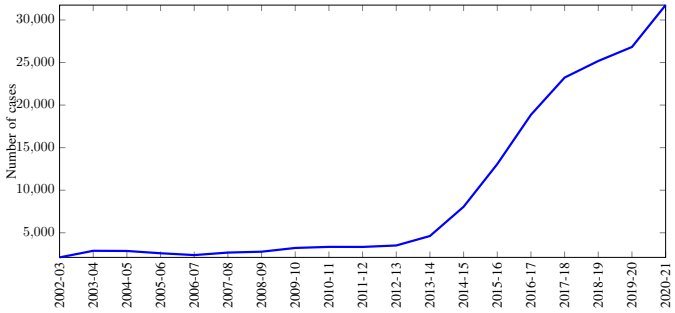

Figure 1: The number of police recorded offenses involving indecent images of children. Reporting years are all from April to March.

The technology, process, and available resource used during forensic investigations have seldom changed in the last 20 years, and due to this restriction, typically only a few concurrent investigations can be handled at any one time,

[1] Crimes categorised under the offence code 86 "Obscene publications, etc and protected sexual material" is not exclusively those of an IIOC nature and it is unfortunately not possible to acquire further breakdown. Data acquired from https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/datasets/crimeinenglandandwalesappendixtables

resulting in a continuously increasing backlog. A reduction in processing time would have a transformative impact, by enabling timely identification of victims, swift intervention with perpetrators to prevent re-offending, and reducing the traumatic psychological effects of any ongoing investigation for the accused and their families. Recent literature presenting a survey of practitioner opinion indicates the future value of more advanced technologies (mainly artificial intelligence) (Sanchez et al. 2019).

The digital investigation process is currently planned and managed manually using finite resourcing. Furthermore, the outcome of the investigation is not known until the end. The investigator is responsible for considering all ongoing cases and determining a viable plan to utilise available resources. The investigator is likely to have competing demands in terms of more cases to process than available resources. In this paper, the process of IIOC investigations is modelled and Automated Planning is applied and evaluated through the use of five case study instances. This paper is structured as follows: the process of an IIOC investigation is presented, before it is modelled. A case study is provided to demonstrate the potential of applying automated deliberation techniques to IIOC investigations through the use of five case studies. Finally, a conclusion is provided where the emphasis is provided on avenues of future research.

Related Work

There is a large body of work demonstrating successful applications of AP in several key cyber security domains, such as penetration testing (Hoffmann 2015), vulnerability assessment (Boddy et al. 2005), red teaming (Yuen 2015), and attack prevention (Khan and Parkinson 2019). Existing literature shows that computer-generated plans can be more efficient than human-created plans as the planner explores a wide array of possibilities in the given domain knowledge (Khan and Parkinson 2017). AP also offers unique benefits to resolve several performance challenges as attack planning is often performed under uncertainty, in large and complex systems, having little to no prior knowledge of how the attack will progress. As discussed before, similar challenges are faced by LEAs in IIOC investigations (Wilson-Kovacs, Rappert, and Redfern 2022; Franqueira et al. 2018), and this work demonstrates that they can be adequately resolved by the novel application of AP techniques. AP not only facilitates in planning and prioritising important cases first, but also optimising the utilisation of already limited law enforcement resources.

Despite the clear need, there is an absence of research undertaken in the area of automated digital forensic processes (Alrumaithi 2018). Some tools have attempted to address shortcomings by providing prioritisation approaches. For example, the Kent Internet Risk Assessment Tool (KIRAT) is based on decision trees that prioritises suspects according to their level of risk of committing contact offences against children (Long et al. 2016), but since publication, there has been no further empirical analysis as to the tool's effectiveness. Another tool, the Child Pornography Offender Risk Tool (CPORT) (Seto and Eke 2015) identifies future directions in case prioritisation by identifying individuals with a high risk of repeat offending prior to arrest; however, it is reliant on manual input from investigators and is inherently limited to handling individuals with prior conviction.

One research study presents a technique (Du and Scanlon 2019) to automatically prioritise suspicious artefacts that might be important to the investigation to reduce the manual effort involved with analysing large volumes of data. The technique is based on supervised machine learning, which leverages file metadata and labelled artefacts encountered in previous cases. The technique only recommends artefacts that are likely to be suspicious and requires further human intervention to generate the final classification. Another similar research (Du 2020) developed a technique where previously known, relevant artefacts are used for training a machine learning model to calculate a similarity score for the unknown, unseen artefacts. The comparison is based on a file's metadata and the approach orders all artefacts by the similarity scores to allow the investigator to focus on the most relevant artefacts first. However, further research is required in this direction to allow case prioritisation ability.

In other recent work, FoRePlan (Katsini et al. 2021), a tool that allows security experts to prepare, manage and execute digital forensics readiness plans to aid detect cyber attacks in Internet-of-Vehicles environment. The tool provides a graphical interface to manually create plans. Based on the attack profile, the tool retrieves the most suitable, relevant plan for data collection, preservation, and analysis actions to be taken. A similar framework, called DF 2.0 (Verma, Govindaraj, and Gupta 2018), takes case information and device/image as input and uses machine learning techniques to efficiently automate the entire investigation process, thereby reducing human workload. However, the framework does not consider any mechanism to prioritise important cases. It is evident that although there has been an absence of exploration into automated digital forensic processes using AP, there are sufficiently mature and sufficient works to warrant its exploration as a viable approach to solve the application challenges.

Problem Statement

Digital forensic investigations are performed in stages by trained experts. A standardised process is followed whereby discrete investigative actions are undertaken. The actions must take place in a specific order and they involve the use of analytical techniques on digital evidence. The typical process of investigating digital resources for IIOC content can be seen as a sequential series of analysis tasks, as illustrated in Figure 2. This process and the estimated times have come through consultation with two UK LEAs. These tasks are:

- **Intelligence**: Processing officer submitting a request to the digital investigation team, including any available intelligence as to the presence of IIOC, whether there is distribution involved, as well as any technology pertinent to the investigation. This is likely to be used for case prioritisation.
- **Data Acquisition**: All available data sources from suspected devices are imaged to create a forensically sound copy for investigation. Different software and hardware

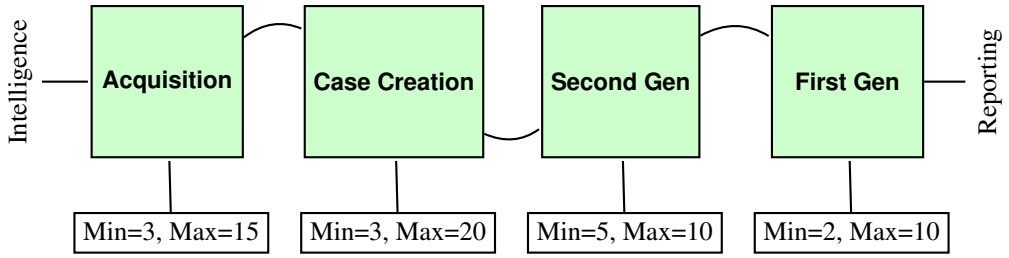

Figure 2: Stages involved in IIOC cases and their estimated minimum and a maximum number of working days required to complete each.

tools are required for acquiring data (referred to as an 'image') from different devices, such as mobile phones, tablets, laptops, and personal computers.

- **Case Creation**: Data acquired is loaded into analysis software, where files are identified for analysis. This includes searching for deleted and recoverable files (referred to as 'carving'). As part of the process, the hash sum for each file will be computed for future verification.
- **Second-generation search**: Known databases (e.g., CAID[2] in the UK) are used to search for matching hash values. This would indicate the presence of a known (i.e., second-generation) image containing IIOC content. This is an automatic process and the images in the known database are assigned a category as to their significance. The presence of the different categories is important from a legal perspective[3].
- **First-generation search**: Images identified during carving that are not matched during second-generation analysis are then examined to see if they contain IIOC content. This involves the use of manual analysis as well as computer vision techniques to identify age, gender, and skin content through pre-trained classifiers. Once discovered, the file hash will be submitted to the hash database for future second-generation matching.
- **Reporting**: Full written report of forensic findings suitable for distribution among legal teams. This is largely a manual process.

Each of the aforementioned tasks is resource-intensive, which is problematic considering the inevitability of finite resources being available for each task. Based on case processing times from the collaborative partners involved in this research, it has been established that the case processing time can range from 13 to 45 shifts. The duration is reported in terms of shifts as although some of the analysis tasks might finish between shifts (e.g., during the night), investigative staff are only there during the working shift and therefore the outcome of the task has to wait for interpretation before any new task is started. Figure 2 provides the breakdown of the anticipated minimum and maximum number of shifts required for each stage. As evident in the figure, all tasks have great variation in duration. The reason for this is that their duration is in relation to the volume of data being analysed.

As with all stages of the investigation, the availability of computing resources is essential to its completion. During the different stages, multiple specialist software solutions are used, with each often focusing on providing specific functionality (e.g., imaging, file carving, etc.). In digital forensics, there are many applications that are used during the investigation for discrete tasks. This is most likely due to the complexity of each application, resulting in it not being viable for commercial organisations to pursue the development of a singular investigative tool. Rather, digital forensic processes often involve the use of different applications in a toolchain that is determined based on the investigative task. The approach to building applications with focused functionality also maximises the use of applications for tasks that are pertinent to all investigations, such as imaging and carving, for example.

Each application is commonly a desktop computer application, utilising all available resources (CPU, memory) during use. The applications are designed to be run on a desktop computer that is not currently being constrained by other executing software, including other instances of the same application. It is also the case that proprietary software is licensed per desktop instance. This results in a sequential ordering of the investigative process, where computers are only able to handle one case at any given time and any findings of the analysis are often only available after the search for second and first-generation IIOC.

One feasible way to increase IIOC investigation in terms of overall case throughput is by increasing the resources available to the investigative team. More specifically, increasing the number of analysis computers and software licenses. However, this alone is not a viable solution. More analysis computers would require more staff to handle the increased number of concurrent cases. Shift patterns could also be changed to maximise available resources to ensure computers are not left idle for lengthy periods. Although increasing available resources would be a welcome addition to any digital forensic unit, a continuous increase would be needed to keep pace with the continuing increase in the number of IIOC-related investigations, as seen in Figure 1. There is therefore a strong need to reconsider current technology and process to make a step change in increasing case throughput.

[2]The Child Abuse Image Database is a database containing hash values of known IIOC

[3]In the UK, thresholds are set by the Crown Prosecution Service as to how many images of each type are required to be prosecuted. This is Cat A) 151 images, Cat B) 1,000 images, and Cat C) 3,000.

Domain Modelling

In this section, the IIOC digital investigation process is modelled to enable the application of Automated Planning (AP). The Planning Domain Definition Language (PDDL) version 2.1 is used as a standardised mechanism to encode the domain, enabling the use of different planning algorithms and the use of the domain and problem instances by researchers within the AP community (Fox and Long 2003). There is a temporal component to the domain model and numerics are required to encode information regarding the size of digital artefacts to investigate and the number of images discovered. The following section describes the encoding of the domain model and problem instances. Only key information is discussed in the interest of brevity and the authors refer readers to the Availability section for information on how the domain model and problem instances used in this paper can be acquired.

Objects, Predicates, and Functions

Based on an understanding of how digital forensic investigations are performed and the resources and exhibits used, it is possible to create an abstract model containing key information, yet remaining sufficiently abstract to not overcomplicate the planning problem. Key object types used in this domain are as follows: a `computer` object is used to represent resources available to the digital investigator for use in investigation and analysis tasks. These are often standard desktop computers. A `case` is used to represent a single investigation case that needs processing.

Predicates are used in the domain model for two primary purposes. The first is through `(computer_available ?computer)` which is used to model if a computer is available to be used for an investigation action. A binary predicate is used for this function as typically the computer used for an investigation can only perform one task at any given time. This is because a task, such as carving, will consume all available resources and any attempt to run multiple at the same time will be detrimental to the overall throughput. A `(case_to_examine ?case)` is used to indicate whether a case still requires further processing before the investigation is complete. This predicate will default to true until an exhaustive analysis of all exhibits has been performed, including the identification and analysis of all identified images.

Numeric fluents (functions in PDDL) are used to store numeric information pertinent to the investigation case. Digital investigations of suspected IIOC content are currently performed sequentially, as presented in Figure 2. This results in 4 key actions per investigation in sequential order, resulting in little opportunity to extract early findings. Furthermore, it presents little need for automated deliberation. However, in the approach presented in this paper, a fundamental shift has been taken to allow each case and investigative action a finite segment of time encoded as a `(time_unit)`, after which information regarding the case (number of files, number of IIOC images, etc.) can be better utilised to prioritise cases. In this work, 60 minutes has been selected as an appropriate time segment.

Each case will be assigned a `(case_priority ?case)` value which is provided as prior intelligence. By default, each case will have a priority of 1 and the value can be increased to establish higher importance. A priority of 0 would indicate that the case has no urgency and, although analysis is required, it is not time-sensitive. For example, this could occur when prosecution thresholds have already been surpassed for a case. Although further processing is important to examine all available images, it is no longer time-sensitive for either prosecution or victim safeguarding. A `(case_processing_time ?case)` fluent is used to record how long a case has been in-process, which will be updated as the investigation is performed.

The imaging process required the following values to be stored in numeric fluents. The first is `(size_disc ?case)` where the size of the storage device being examined is recorded. The fluent `(size_imaged ?case)` is used to store how much of the physical device has been imaged, and the value `(image_speed_min)` is used to represent the speed of the imaging process, in MB per min. In terms of the carving process, `(size_carved ?case)` is used to store how much data has been carved during the carving process. The carving speed is known beforehand and is represented in the in `(carving_speed_min)`, where the carving speed in MB per minute is stored. These performance values can be changed depending on available hardware.

As carving is taking place, files will be discovered that need investigation. The number of files will be held in two separate variables of `(first_gen_files_to_proc ?case)` and `(second_gen_files_to_proc ?case)`. The reason for storing this value in two separate variables is so that separate processes of first and second-generation search can decrement the value as they are examined. In the same way as carving, the speed of first and second generation analysis in terms of files per minutes available and held in `(first_gen_speed_min)` and `(second_gen_speed_min)`. As the first and second-generation analysis techniques are being performed, the number of images discovered as valid first and second-generation IIOC images are stored in (first_gen_speed_min) and `(second_gen_speed_min)`, respectively. A `(discovery)` fluent is used to keep track of the detection of IIOC images and at what point in the investigation they are discovered. A lower value means that the images are discovered earlier in the investigation.

Finally, three values are used to store the number of images required for each different prosecution threshold. For example, `(threshold_a)` is used to store the threshold value of 151.

Operators

In the produced domain model, actions have been created to model (1) creating a forensically sound image of a device, (2) carving the image to identify and recover all files, (3) analysing the files for second-generation IIOC content, and (4) analysing the files for first-generation content. It is important to note that this is an iterative planning and execution problem, whereby as information is discovered during

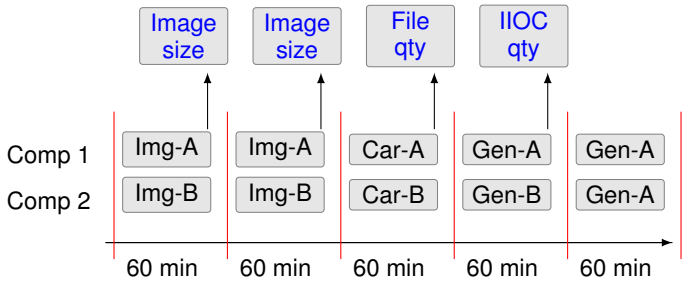

Figure 3: Example timeline illustrating two cases (A, B) being processed concurrently in 60-minute segments, Img = imaging, car = carving, gen = first and second generation searching.

the investigation, planning is performed once again to ensure cases of the highest priority, or those where IIOC images are being discovered, are prioritised until prosecution thresholds are reached. The following subsections describe these action types. Figure 3 provides a timeline illustration for processing two cases where two computers are available. The two cases follow a similar path, whereby the first two 60 minute segments are occupied by imaging activities, followed by carving for 60 minutes to identify all files, and then analysis tasks seeking to find first and second-generation images. Note that this illustration is a simplification and that these activities have been merged into one 'Gen' action for illustration purposes. Although the actions are the same in the first four 60-minute segments, the fifth segment is different as case 'A' is being processed on both computers simultaneously. The reason for this is due to the information output after each 60 minute segment that feeds into the next planning round. This information is dependent on the action execution. For example, imaging outputs how much data has been imaged thus far, whereas carving outputs how many files have been discovered. IIOC analysis processes will report on how many positive matches they have found, which is used to prioritise cases as seen in the figure.

Imaging This action is used to encode the imaging process, whereby digital evidence is processed to create a forensically sound byte image file. The entire device needs to be imaged before any further processing can take place. As seen in Figure 4, the imaging encoding is a temporal action whereby providing there is still disc space left to image, a computer is available, and the case is still active, the imaging process will be scheduled to take place for a `(time-unit)`. The encoding models how much data will be processed during that time using provided information. Note that although the imaging process has to be complete before other actions can be considered, adopting the same time segment approach enables the planner to prioritise the use of the resources for other cases.

Carving The carving action is used to model the process of taking the image file and starting to process it to extract files. In addition to extracting files present in the underlying file system, the carver will also identify and piece together file fragments, belonging to deleted files. The carving action has conditions to check that the imaging process has finished, as well as checks that there is still data available to be carved. Similarly to the imaging action, the carving action estimates how much carving will take place in the 60-minute segment by using predetermined values on how much data the carver can handle per minute.

```
(:durative-action imaging
   :parameters(?PC - computer ?C - case)
   :duration(= ?duration (time_unit))
   :condition
   (and
    (at start (computer_available ?PC))
    (at start (case_to_examine ?C))
    (at start (< (size_imaged ?C)
      (size_disc ?C))))
   :effect
   (and
    (at start(not
      (computer_available ?PC)))
    (at end(computer_available ?PC))
    (at end(increase(size_imaged ?C)
      (*(time_unit)(image_speed_min))))
    (at end(increase
      (case_processing_time ?C)
      (time_unit)))
   )
)
```

Figure 4: The PDDL encoding of the imaging action, which models the process of imaging a digital resource.

First and Second-Generation Analysis Searching for first and second-generation images in identified files is encoded in two actions. Although they are almost identical in encoding, it is necessary that they are separate as they are performing distinctly different analysis tasks. The process and software involved with searching for first-generation are different from second-generation. Figure 5 provides the encoding of the action for second-generation analysis. The action has a precondition to ensure that there are images in need of analysis, and as with other actions, the number is decreased by the capability of the computing resources. This action also has a `(discovery)` fluent which is increased by the number of images discovered multiplied by the minutes that the investigation has been running for and the case priority. This provides a measurement of how significant the discovery has been and will be used for prioritisation.

Case Complete & Threshold Reached

Finally, the domain model contains a series of non-durative 'housekeeping' actions that are used to either state that the investigation has been completed (i.e., all imaging and analysis tasks have been performed), or to state that prosecution thresholds have been reached and the case priority is reduced.

```
(:durative-action secondgen
   :parameters(?PC - computer ?C - case)
   :duration(= ?duration (time_unit))
   :condition
   (and
    (at start (computer_available ?PC))
    (at start (>
      (second_gen_files_to_proc ?C)0))
    (at start (case_to_examine ?C))
   )
   :effect
   (and
    (at start(not(computer_available ?PC)) )
    (at end(computer_available ?PC))
    (at end(decrease
      (second_gen_files_to_proc ?C)
      (*(time_unit)(second_gen_speed_min))))
    (at end(increase(discovery)
      (*(+(case_processing_time ?C)
      (second_gen_discovered ?C))
      (case_priority ?C))))
    (at end(increase
      (case_processing_time ?C)
      (time_unit))))
   )
)
```

Figure 5: The PDDL encoding of the second-generation analysis action

Initial and Goal State

The initial state specifies information to describe currently available resources and cases to process. For example, `(computer_available PC1)` states that PC1 is available. In terms of technical details and capabilities of the computers, `(=(size_disc case1)102400)` states that the case includes a 100GB hard drive and `(=(image_speed_min)100)` states that the imaging speed is 100MB per minute.

In addition to information encoded to represent static information regarding the case, fluents are used to encode information generated during carving, first and second-generation analyses. This includes `(=first_gen_files_to_proc ?case)0)` which is used to specify how many files have been discovered during the carving process and can be used for first-generation analysis. Similarly, `(=(firstd_gen_discovered ?C)0)` is used to denote how many first-generation IIOC images have been discovered. The same set of values exists for the second-generation analysis action. Both these values are updated after each 60-minute interval as the unknown and determined at run-time.

Plan Metric

The following plan metric is introduced for optimising the quality of generated plans:

- `(:metric minimize (discovery))`

As the objective of this research is to prioritise the handling of cases containing IIOC content, it is necessary to introduce a `minimize` metric to priorise further analysis of cases where IIOC has been discovered as soon after discovery as possible. More specifically, as the `(case_processing_time ?case)` fluent increases, the impact on the `(discovery)` metric is greater.

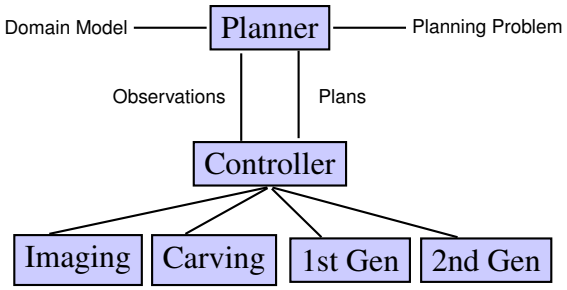

Figure 6: Illustration of the planning and execution system.

Plan Generation and Execution

In this application, it is necessary to continue to monitor the execution of the plan and re-plan after every execution segment (60 minutes), taking into consideration information on files discovered during carving and images discovered to be suspecting of IIOC. An illustration of the application framework is presented in Figure 6. The LPG-td planner (Gerevini, Saetti, and Serina 2003) is selected due to its exploitability in real-world planning applications, its good support of PDDL features, and its good performance. The domain model and planning problem is provided to the planning tool, which is overseen by the controller. In this instance, the controller is a C# application that controls the execution of digital forensic tools. The controller executes the plan and, after each execution segment, will update the problem instance before re-planning. The planner is allowed 1-minute CPU time to generate plans before the one producing the lowest metric is selected. The tools used in this example are *FTKImager* for creating a forensically sound copy of the storage device and the *Bluebear carver* is used for file carving. In terms of second-generation matching, the standardised hash-matching approach is performed using hash lists provided by LEAs, such as CAID in the UK. The detection of first-generation content requires the use of pre-trained classifiers to provide probability measures on images relating to age, gender, and skin coverage, among others.

Case Study

In this work, five case study benchmark instances are used to demonstrate the potential of this approach. The five benchmark instances are constructed in collaboration with LEAs to ensure they are representative of current investigations. It is worth emphasizing that only benchmark instances can be constructed and used for this research that do not contain IIOC content. We add images to discover by following two principles: 1) we introduce hash values of images into the hash list used in second-generation search, and 2) we add images that would result in positive matches in terms of age (<18 years old) and use these to represent images to discover. These images are taken from publicly available

	1	2	3	4	5
Size	100GB	200GB	150GB	1TB	2TB
# Images	50,000	60,000	70,000	34,000	20,000
# 1st gen	0	0	1000	100	0
# 2nd gen	0	0	0	2000	0
Imaging	0.7	1.4	1.1	7.1	14.2
Carving	1.0	2.0	1.5	10.2	20.3
2nd gen	0.6	0.7	0.8	0.4	0.2
1st gen	3.5	4.2	4.9	2.4	1.4
Total	5.8	8.3	8.3	20	36.2

Table 1: Details of case study cases used for comparative analysis. The top provides specifics of the case and the bottom half provides how many days are required for that part of the processing. In terms of the processing capabilities of the computers used, imaging was 100MB/min, carving 70MB/min, second-generation search 60 files/min, and first-generation 10 files/min

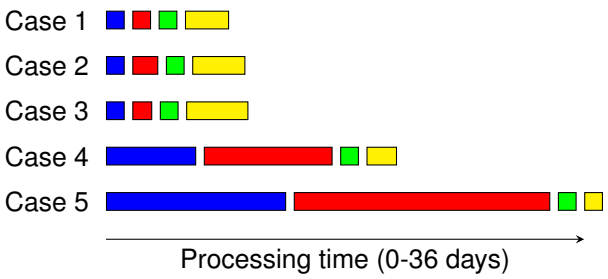

Figure 7: Timeline of five case study cases. Each square represents an activity and is the duration (length). Blue is for acquisition, red is for carving, green is for second-generation analysis and yellow is for first-generation analysis.

data sets, such as the UTK Datasets (Zhang and Qi 2017). Table 1 presents the information from the five benchmark instances. The top half of the table. The cases are of varying sizes, numbers of images, and also a number of IIOC images to detect. In terms of duration, the cases require between 5.8 to 36.2 days to process using the traditional approach. In addition, from Table 1 it is evident that only two (3 and 4) of the five cases have IIOC to discover. Case 3 has 1,000 first-generation images to discover, whereas case 4 has 100 first-generation and 2,000 second-generation. An important observation here is that the 100 first-generation images in case 4 are not sufficient to reach a prosecution threshold and the second-generation images will need to be discovered.

Figure 7 provides an illustration of the five cases where coloured squares are used to denote each investigative action and their length. Based on an understanding of the current process, information regarding the discovery of IIOC content is only known at the end of the entire process, so for cases 3 and 4, that is after 8.3 and 20 days, respectively. Furthermore, the five cases are occupying an investigative computer each and are segmenting any other cases being processed.

The same five cases are then handled by the application presented in this paper. As previously mentioned, the approach presented in this paper plans in time segments of 60

Case	Completion Time (days)	Time Difference	IIOC Discovery Time	Time Difference
1	5.9	+2%	n/a	n/a
2	8.8	+6%	n/a	n/a
3	8.4	+1%	7.1	14%
4	13.3	-33%	12.4	38%
5	22.6	-38%	n/a	

Table 2: Results from applying automated planning to the five case study examples

minutes. All problem instances were initially encoded using the information provided in Table 1. The C# application (the 'Controller' in Figure 6) generates a plan (using LPG-td) before executing the 60 minute segment for each case. Following the execution, the problem instance is then updated to account for information on whether the image is ready for carving, whether carved files have been identified, and whether first or second-generation IIOC content has been discovered. The analysis of all cases has been done till completion, as shown in Table 2. Further, it is evident that although the first three cases have been completed requiring a slightly longer duration, those of a larger size (cases 4 and 5) have been completed with a significant reduction of 33% and 38%, respectively. This is because once the first three cases have been completed, the extra resources become available for the remaining two cases. This results in many analysis tasks being planned in parallel and results in them finishing much more quickly. In a live investigative setting, once the first three cases are complete, it is likely that three more cases from the backlog would be selected to be processed. Although the use of automated planning would not result in faster throughput if the number of cases being processed is large, it does allow the parallelisation of tasks should an investigation require a more timely conclusion. In addition, it also allows the technique to handle more cases than computers available, which is something not currently possible. Although this would be detrimental to the processing duration of all cases being concurrently handled, a case would be priorised and allocated more resources should a case be discovered to contain IIOC content.

Table 2 also provides information on when IIOC content was discovered within a case and when prosecution thresholds were reached. This is significant, as the current way of working does not produce any insights or outcomes back until the entire investigation process is complete. When cross-referencing the 'IIOC Discovery Time' from Table 2 with the 'Total' time in Table 1, it is evident that there is a good reduction in the duration taken to inform the investigative team of IIOC content. This is a 14% (1.2 hours) and 38% (7.6 hours) reduction for cases 3 and 4, respectively. This is significant as any reduction in IIOC discovery time can result in more timely prosecution and victim safeguarding.

Figure 8 provides a graphical illustration of an excerpt from the plan. Each case is represented by a series of executed actions that are colour coded to match their activity type. Each segment represents 60 minutes of execution and

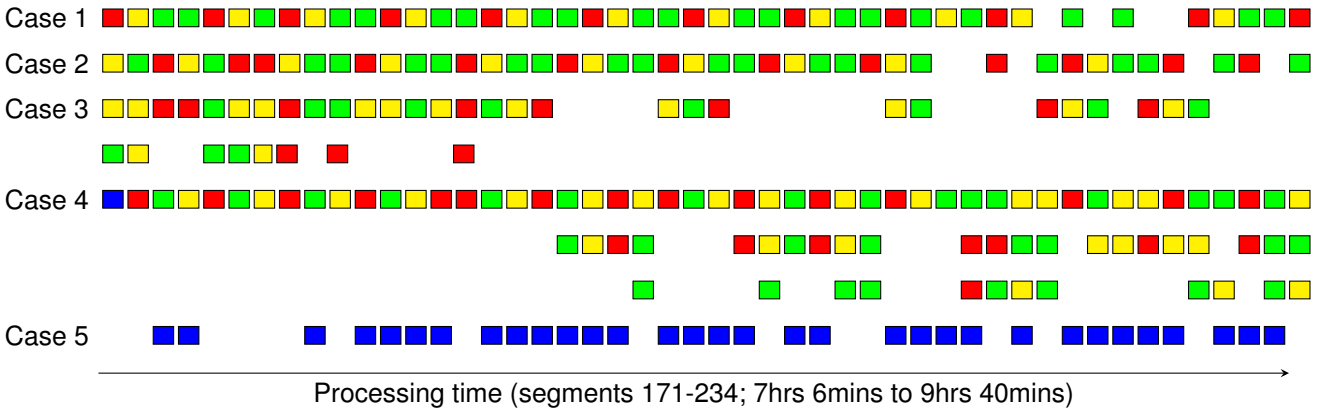

Figure 8: plan excerpt illustrating the planned and executed plan over 48 hours. The illustration shows the planned activities for each 60-minute segment. Blue is for acquisition, red is for carving, green is for second-generation analysis and yellow is for first-generation analysis

the excerpt covers the period of activity from 7hrs 6mins to 9hrs 40mins. This is region 171-234 in terms of 60-minute segments. From the excerpt, it is immediately evident that cases 1-3 are in a cycle of performing carving actions before subsequent first and second-generation analysis actions. This is because, in a 60 minute segment of carving, new files are discovered that can be analysed for IIOC content. The planner is selecting these actions once images are available based on the potential to minimize the reward function. It is also immediately evident that case number 5 is still in the imaging phase of the analysis. As previously specified in the modelling section, carving can only take place once imaging has finished. This is by no means a surprise for case 5 as it is the one with the largest size to process (2TB). Although no analysis of case evidence is taking place during the imaging phase, it is evident that the planner may choose to prioritise resources to other cases where deemed appropriate.

Cases 3 and 4 demonstrate how actions are planned based on outcomes from previous analysis actions. For example, for case 3, it is evident that the planner is allocating more resources to the case, prioritising it over case 5. This is because IIOC content is being discovered during first and second-generation analysis, which results in the planner choosing to allocate more analysis actions to the case. This takes place in the first third of the excerpt before it stops and the planned actions for case 3 significantly reduce. This is because the prosecution threshold for case 3 has been reached, and therefore priority is allocated to other cases. This is a significant contribution as it makes more resources available to other cases once sufficient information has been discovered for the investigating officers to prosecute. However, it should be noted that the case continues to be investigated until complete, which is required in all investigations. The excerpt also demonstrates that case 4 completes its last imaging action and carving can take place. It is also evident that after the first third of the excerpt, and case 3 passed the prosecution threshold, the planner started to increase planned actions for case 4. This is because IIOC content has started to be discovered and the planner is prioritising analysis actions. As the contents of these cases are known beforehand, this behaviour makes sense as it is known that cases 1, 2, and 5 do not contain IIOC content. The excerpt presented in Figure 8 and the values presented in Table 2 demonstrate that the approach is intelligently planning and scheduling investigative actions based on available resources and insights generated by previous actions.

Availability

Domain models and problem instances used in this paper are available from the authors upon request.

Conclusion

In this paper, the process of digital forensic investigations of cases containing suspected IIOC contents is presented. The scale of the challenge is outpacing resource capability and it is evident that innovation is needed to improve investigative processes. The use of automated planning is explored to better allocate resources, and to utilise and report early findings during the investigation, overcoming the current way of working whereby information on IIOC content is only reported at the end of the investigation. This delay has significant consequences for safeguarding and timely prosecution. In this wor, the process is modelled following the same investigative actions currently performed; however, a significant change is that each investigative action is performed in a 60-minute segment before re-planning is performed based on monitoring execution. This enables early discoveries and their use in prioritsation, as well as performing concurrent actions. Empirical analysis of 5 case studies resulted in a reduction of 36% in processing time and a 26% reduction in time required to discover IIOC content. In future work, further experimentation is to be performed to examine the approach's capability when handling a larger number of cases, as well as investigate techniques to increase efficiency. In addition to the analysis using simulated benchmarks, trials within LEAs will enable access to real historic cases.

References

Alrumaithi, A. M. 2018. *Prioritisation in Digital Forensics: A Case Study of Abu Dhabi Police*. Liverpool John Moores University (United Kingdom).

Ayers, D. 2009. A second generation computer forensic analysis system. *digital investigation*, 6: S34–S42.

Boddy, M. S.; Gohde, J.; Haigh, T.; and Harp, S. A. 2005. Course of Action Generation for Cyber Security Using Classical Planning. In *ICAPS*, 12–21.

Du, X. 2020. Alleviating the Digital Forensic Backlog: A Methodology for Automated Digital Evidence Processing. *School of Computer Science, University College Dublin*.

Du, X.; and Scanlon, M. 2019. Methodology for the automated metadata-based classification of incriminating digital forensic artefacts. In *Proceedings of the 14th International Conference on Availability, Reliability and Security*, 1–8.

Fox, M.; and Long, D. 2003. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of artificial intelligence research*, 20: 61–124.

Franqueira, V. N.; Bryce, J.; Al Mutawa, N.; and Marrington, A. 2018. Investigation of Indecent Images of Children cases: Challenges and suggestions collected from the trenches. *Digital Investigation*, 24: 95–105.

Garfinkel, S. L. 2010. Digital forensics research: The next 10 years. *digital investigation*, 7: S64–S73.

Gerevini, A.; Saetti, A.; and Serina, I. 2003. Planning through stochastic local search and temporal action graphs in LPG. *Journal of Artificial Intelligence Research*, 20: 239–290.

Hoffmann, J. 2015. Simulated penetration testing: From" dijkstra" to" turing test++". In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 25, 364–372.

Katsini, C.; Raptis, G. E.; Alexakos, C.; and Serpanos, D. 2021. FoRePlan: Supporting Digital Forensics Readiness Planning for Internet of Vehicles. In *25th Pan-Hellenic Conference on Informatics*, 369–374.

Khan, S.; and Parkinson, S. 2017. Towards Automated Vulnerability Assessment.

Khan, S.; and Parkinson, S. 2019. Discovering and utilising expert knowledge from security event logs. *Journal of Information Security and Applications*, 48: 102375.

Long, M.; Alison, L.; Tejeiro, R.; Hendricks, E.; and Giles, S. 2016. KIRAT: Law enforcement's prioritization tool for investigating indecent image offenders. *Psychology, Public Policy, and Law*, 22(1): 12.

Sanchez, L.; Grajeda, C.; Baggili, I.; and Hall, C. 2019. A Practitioner Survey Exploring the Value of Forensic Tools, AI, Filtering, Safer Presentation for Investigating Child Sexual Abuse Material (CSAM). *Digital Investigation*, 29: S124–S142.

Seto, M. C.; and Eke, A. W. 2015. Predicting recidivism among adult male child pornography offenders: Development of the Child Pornography Offender Risk Tool (CPORT). *Law and human behavior*, 39(4): 416.

Verma, R.; Govindaraj, J.; and Gupta, G. 2018. DF 2.0: Designing an automated, privacy preserving, and efficient digital forensic framework.

Wilson-Kovacs, D.; Rappert, B.; and Redfern, L. 2022. Dirty Work? Policing Online Indecency in Digital Forensics. *The British Journal of Criminology*, 62(1): 106–123.

Winder, B.; Blagden, N.; Armitage, R.; Duncan, K.; Wakeham, A.; Roberts, L.; and Berti, C. 2020. The experiences of non-offending partners of individuals who have committed sexual offences. Recommendations for practitioners and stakeholders.

Yuen, J. 2015. Automated cyber red teaming.

Zhang, S. Y., Zhifei; and Qi, H. 2017. Age Progression/Regression by Conditional Adversarial Autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.