

Formalising Event Time Bounding in Digital Investigations

Pavel Gladyshev, Ahmed Patel
School of Computer Science and Informatics
University College Dublin, Ireland

Abstract

A timestamp is a clock reading attached to a unit of data. Timestamps are widely used in computing and seem to offer an easy way to determine the time of events in digital investigations. Unfortunately, the ability of users to change clock settings reduces the evidential weight of timestamps. Alternative methods of estimating times of events are often needed to corroborate timestamps. One such method is to “sandwich” the unknown time of an event between known times of causally connected events. For example, if event *A* caused event *B* and event *B* caused event *C*, then the time of *B* must be between the times of *A* and *C*. This type of reasoning is sometimes called “event time bounding.” This paper defines event time bounding as a mathematical problem and presents an algorithm for solving it ¹.

Section 1: Introduction

The determination of event times is an important and difficult task in digital forensics as demonstrated by the following example, which was originally described in Bates (1997):

A managing director of some company, Mr. C was blackmailed. During the investigation, a computer belonging to Mr. A, a friend of Mr. C, was subjected to forensic analysis. Several deleted fragments of the blackmail letter were discovered. They provided a recognised sequence of changes to the blackmail letter over a period of time. Mr. A claimed that the letter was created by somebody else in his absence during his holidays. Although some of the deleted fragments had timestamps associated with them, those timestamps were considered to be unreliable, because they indicated the setting of the internal computer clock which may have had no relevance to real world dates and times.

To support the timestamp evidence, the investigators attempted to corroborate the guilt of Mr. A by the following reasoning:

1. A fragment of the blackmail letter was found in the slack space of a file that happened to be a letter to a third party (say, a bank manager).
2. The bank manager confirmed receiving that letter on the date before the beginning of Mr. A's holidays.
3. The investigators reasoned that, because of the way the slack space is formed (see Chapter 9 of Casey, 2004 for details), the blackmail letter must have been saved to the disk *before* the second letter was created, posted and received by the bank manager (see Figure 1 below).

¹ The research reported in this article was sponsored by Teltec Ireland in 1999-2000.

When Mr. A was confronted with this reasoning, he admitted writing the blackmail letter and pleaded guilty to the charges of blackmail.




Figure 1. Ordering of events in the blackmail incident

The example demonstrates how causal connections between events can be used to determine possible time intervals for events whose time is unknown. This kind of reasoning is sometimes called event time bounding. The rest of this paper shows how this kind of reasoning can be formalised and automated. Section 2 defines a mathematical model of causally connected events. Section 3 then describes a method for calculating possible time intervals of events using the defined model. This is followed by an application of the developed method to a fictional example in Section 4. The conclusions for the paper are then given in Section 5.

Section 2: Event times and causal connections between events

For the purposes of this paper, events are considered to be instantaneous. The *time of an event* is defined as the moment when the statement describing the event becomes true in the world. For example "the meeting started at 17:00," "the meeting ended at 18:00," "a letter was posted at 12:37," etc.

Hereinafter, events are denoted by capital letters A, B, C, and E with or without subscripts: e.g. A , B , C , C_i , E , E_j . The time of events is denoted T^A , T^B , T^C , T^{C_i} , T^E , and T^{E_j} respectively.

In the real world, the time cannot be measured exactly. The statement, "a letter was posted at 12:37," effectively means that the letter was posted some time after clock reading became 12:37 but before it became 12:38. Thus, for any real world event E there are the earliest and the latest times between which the event happened. These times are denoted as T_{\min}^E and T_{\max}^E respectively. By definition $T_{\min}^E \leq T^E \leq T_{\max}^E$.

Representation of time as an interval allows different degrees of imprecision. For example, the time 12:50 can be viewed as $T_{\min}^E = 12:50$ and $T_{\max}^E = 12:51$. Similarly, the time "around 5 pm" can be viewed, for example, as $T_{\min}^E = 16:00$ and $T_{\max}^E = 18:00$. The time, which is totally unknown, can be represented as $T_{\min}^E = -\infty$ and $T_{\max}^E = +\infty$.

Thus, the task of finding the exact time of event E can be replaced by the task of minimising the difference $\Delta T^E = T_{\max}^E - T_{\min}^E$. We will use causal connections between events to minimise ΔT^E .

To reason about causal connections between events, a suitable mathematical representation of causality is needed. One such representation is the *happened-before* relation, which was introduced by Leslie Lamport (1978) as a tool for reasoning about time in distributed systems. Applied to the forensic context, the *happened-before* relation can be defined as follows.

An event A *happened-before* event B if either (or both) of the following is true:

- B uses results of A . For example, before Mr. X posted the letter to the bank manager (event B), he first printed that letter (event A).
- A precedes B in the usual course of business of some organisation or during normal operation of a machine. For example, Mr. X posted letter (event A) before the post office closed for the night (event B).

Hereinafter the *happened-before* relation is denoted by an arrow: $A \rightarrow B$.

For some event E , there may be several events that *happened-before* E , and E itself may be a *happened-before* event for several other events (see Figure 2).




Figure 2. Happened-before relation

The *happened-before* relation possesses the following properties:

- It is *transitive*: if A happened before B and B happened before C , then A happened before C :

$$(A \rightarrow B) \wedge (B \rightarrow C) \Rightarrow A \rightarrow C$$

- It is *non-reflexive*: A cannot happen before itself.

This property also means that circular *happened-before* relations are disallowed, because $(A \rightarrow B) \wedge (B \rightarrow C) \wedge \dots \wedge (X \rightarrow A) \Rightarrow A \rightarrow A$

- It is *partial*: if two events are not causally connected, we cannot say that one *happened-before* another.

A group of events together with happened-before relations between them can be represented as a *directed acyclic graph* or *dag*. An example of such a representation is given in Figure 3 below. Events are represented by vertices (circles). The *happened-before* relations between pairs of events are represented by arcs (arrows).




Figure 3. A group of events represented as a dag

Section 3: Calculating time limits of causally connected events

As its name states, the temporal property of the *happened-before* relation is that if event *A* *happened-before* event *B*, then the time of event *A* is less than the time of event *B*:

$$A \rightarrow B \Rightarrow T^A < T^B$$

Because of the transitivity of the *happened-before* relation, the same is true for any event that *happened-before* *B*:

$$(Y \rightarrow X) \wedge \dots \wedge (C \rightarrow B) \wedge (B \rightarrow A) \Rightarrow Y \rightarrow A \Rightarrow T^Y < T^A$$

If event *A* *happened-before* event *B* and event *B* *happened-before* event *C*, then the time of event *B* is bounded by the times of events *A* and *C*:

$$(A \rightarrow B) \wedge (B \rightarrow C) \Rightarrow T^A < T^B < T^C \tag{1}$$

This property can be used to minimise the time interval $\Delta T^B = T_{\max}^B - T_{\min}^B$ of event *B*.

Suppose that we do not know the exact time T^B but have some working estimates T_{\min}^B and T_{\max}^B . Now, in the course of investigation we discover the exact time T^A of the event *A*. Because of (1), the exact time T^B is greater than T^A . Thus, if our current estimate $T_{\min}^B < T^A$ we can improve it by choosing new $T_{\min}^B = T^A$.

Similarly, if we discover the exact time T^C of event *C*, such that $T_{\max}^B > T^C$, then we can improve our working estimate T_{\max}^B by choosing new $T_{\max}^B = T^C$.

If there are several events A_1, A_2, \dots, A_N that *happened-before* *B* and several events C_1, C_2, \dots, C_M , such that *B* *happened-before* any of them, we can calculate the best estimates T_{\min}^B and T_{\max}^B by choosing

$$\begin{aligned} \text{New } T_{\min}^B &= \max(T_{\min}^B, T^{A_1}, T^{A_2}, \dots, T^{A_N}) \text{ and} \\ \text{New } T_{\max}^B &= \min(T_{\max}^B, T^{C_1}, T^{C_2}, \dots, T^{C_M}) \end{aligned} \tag{2}$$

A graphical interpretation of the above formulae is given in Figure 4 below.

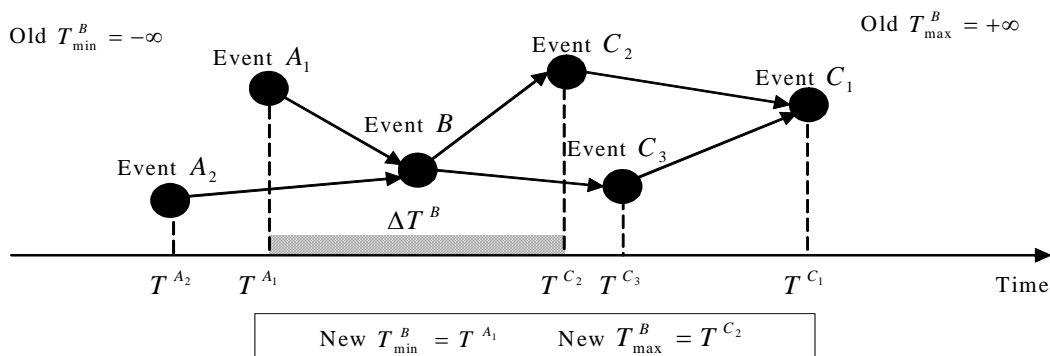


Figure 4. Choosing the best estimates for T_{min}^B and T_{max}^B according to formulae (2)

As was mentioned in Section 0, the time of events in the real world cannot be measured exactly. It means that in place of T^A and T^C in formulae (2) we have to use some value between T_{min}^A and T_{max}^A and some value between T_{min}^C and T_{max}^C . If we ignore delays between events A, B and C , then to preserve the meaning of T_{min}^B and T_{max}^B we have to use T_{min}^A in place of T^A and T_{max}^C in place of T^C (a proof of this claim is given in Appendix 1) as shown in Figure 5.




Figure 5. Time interval of event B restricted by time intervals of events A and C

To reflect this, equations (2) can be rewritten as follows:

$$\begin{aligned}
 \text{New } T_{min}^B &= \max(T_{min}^B, T_{min}^{A_1}, T_{min}^{A_2}, \dots, T_{min}^{A_N}) \quad \text{and} \\
 \text{New } T_{max}^B &= \min(T_{max}^B, T_{max}^{C_1}, T_{max}^{C_2}, \dots, T_{max}^{C_M})
 \end{aligned}
 \tag{3}$$

Formulae (3) ignore delays between causally connected events. However, these delays can be used to further improve the above formulae.

Between two causally connected events in the real world, there always exists a minimal delay imposed by the speed of light. For particular types of events, there are similar but larger limits caused by other reasons. For example, a man cannot travel from place to place faster than 2000 miles per hour (unless he is a superman or an astronaut).

To take advantage of minimal delays, we specify a minimal delay d for each *happened-before* relation, and denote it as follows:

$$A \xrightarrow{d} B$$

This affects the temporal property of the *happened-before* relation in the following way:

$$A \xrightarrow{d} B \Rightarrow T^A + d < T^B$$

In other words, event *B* cannot happen earlier than the time of event *A* plus the minimal delay between *A* and *B*.

Because of the transitivity of the *happened-before* relation, if two events are causally connected by a *single* chain of events, the minimal delay between them can be calculated as a sum of minimal delays along the chain:

$$\begin{aligned} (A \xrightarrow{d_1} B) \wedge (B \xrightarrow{d_2} C) \wedge \dots \wedge (X \xrightarrow{d_n} Y) &\Rightarrow \\ T^A + d_1 + d_2 + \dots + d_n < T^Y &\Rightarrow \\ A \xrightarrow{d_1 + d_2 + \dots + d_n} Y & \end{aligned}$$

If two events are connected by several parallel chains of events, the delay *largest across all chains* should be used.

By the same line of reasoning that led to (2) and (3), note that

$$(A \xrightarrow{d_{AB}} B) \wedge (B \xrightarrow{d_{BC}} C) \Rightarrow (T^A + d_{AB}) < T^B < (T^C - d_{BC})$$

which leads to the revised formulae:

$$\begin{aligned} \text{New } T_{\min}^B &= \max [T_{\min}^B, (T_{\min}^{A_1} + d_{A_1B}), (T_{\min}^{A_2} + d_{A_2B}), \dots, (T_{\min}^{A_N} + d_{A_NB})] \\ &\text{and} \\ \text{New } T_{\max}^B &= \min [T_{\max}^B, (T_{\max}^{C_1} - d_{BC_1}), (T_{\max}^{C_2} - d_{BC_2}), \dots, (T_{\max}^{C_M} - d_{BC_M})] \end{aligned} \tag{4}$$

Figure 6 shows how minimal delays affect calculation of the time interval for the event *B*.

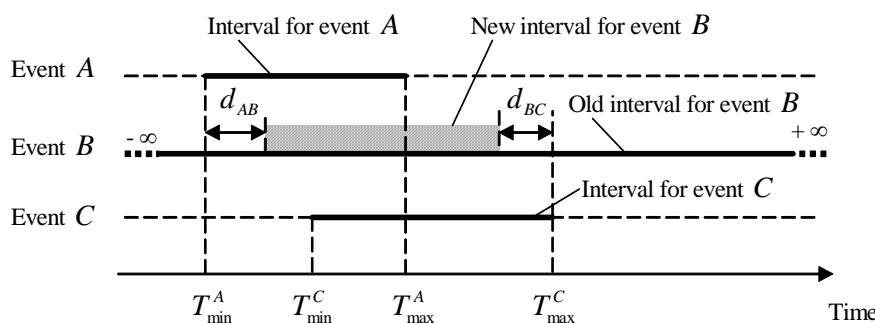


Figure 6. Time interval of event *B* restricted by time intervals of events *A* and *C* with minimal delays d_{AB} and d_{BC}

It must be pointed out, however, that incorrectly specified delays may lead to incorrect results. Suppose that the actual minimal delay d on some *happened-before* relation $A \xrightarrow{d'} B$ is smaller than the specified delay d' . If $T_{\min}^A + d'$ is chosen as the new value for T_{\min}^B according to formulae (4), then since $d < d'$ it is true that

$T_{\min}^A + d < T_{\min}^B$. Hence, the actual T^B may be such that $T_{\min}^A + d < T^B < T_{\min}^B$. In other words, event B may have happened outside the calculated time interval! Unless the value of the minimal delay can be determined with absolute certainty, it is safer to specify the minimal delay equal 0.

Now the high-level algorithm for our method can be specified.

A method for calculating time limits of causally connected events:

1. Identify events
2. Specify $T_{\min}^{E_i}$ and $T_{\max}^{E_i}$ for each event, assuming $T_{\min}^{E_i} = -\infty$ and $T_{\max}^{E_i} = +\infty$ for events whose time is unknown.
3. Identify and motivate *happened-before* relations between events.
4. Identify and motivate minimal delay for every identified *happened-before* relation. Assume zero delay, if no particular time limit can be arguably specified.
5. For each event E_i find *all* events that *happened-before* E_i . It includes events connected to E_i through a chain of events. Similarly, find *all* events for which E_i is a *happened-before* event.
6. For each pair of events E_i and E_j , such that E_i *happened-before* E_j , calculate minimal delay $d_{E_i E_j}$ as a sum of minimal delays along the chain of events connecting E_i and E_j . If E_i and E_j are connected by several parallel chains of events, choose the delay *largest across all chains*.
7. For each event E_i calculate $T_{\min}^{E_i}$ and $T_{\max}^{E_i}$ according to formulae (4).

Note that while steps 1 to 4 of the above method are situation specific, steps 5 to 7 are generic and can be implemented as a computer program. An algorithm implementing steps 5 to 7 is given in Appendix 2.

Section 4: Example application of the method

The method is based on simple observations about the nature of events and their causal connections. It can be easily understood and used by interested parties, such as IT administrators, intrusion analysts, police investigators, etc.

Note that the correctness of the time limit calculations depends on the correctness of the event graph that describes the situation in question. Special care must be taken not to introduce non-existing causal connections and minimal delays. Additional investigation may be required before some of the events or connections may be added.

To see how the described method can be used in practice, consider the following example situation:

During a fraud investigation, employees X, Y, and Z of some company were interviewed. Employee X said that on the day D he entered the company building in the morning. During the day, X logged into the company computer network, visited the canteen (in the same building) and wrote a reply to an e-

mail from the employee Y. At the end of the day, employee X logged out from the computer network and left the company building.

A security guard confirmed that employee X entered the building at 08:30 and left the building at 17:30. He was sure that X could not enter or leave the building except through the main entrance. Employee Y said that she sent an e-mail to employee X around noon. Finally, employee Z said that he saw X in the canteen after 15:00.

The first step is to identify the events. Table 1 lists events, their descriptions and time intervals as identified by the text.

Table 1. Events

<i>Event</i>	<i>Description</i>	T_{\min}	T_{\max}
E_1	X entered the building	08:30	08:31
E_2	X logged into the company computer network	$-\infty$	$+\infty$
E_3	X entered the company canteen	$-\infty$	$+\infty$
E_4	X left the company canteen	$-\infty$	$+\infty$
E_5	X sent an e-mail to Y	$-\infty$	$+\infty$
E_6	X logged out of the company computer network	$-\infty$	$+\infty$
E_7	X left the building	17:30	17:31
E_8	Y sent an e-mail to X	11:00	13:00
E_9	Z saw X in the canteen	15:00	$+\infty$

The second step is to identify causal connections between events. Items 1 to 5 in the list below define connections between events whose timing is known. Items 6 and 7 define connections that ought to exist if the events that happened according to X's story did happen.

- 1) Z could not have seen X in the canteen before X entered the canteen: $E_3 \rightarrow E_9$
- 2) Before X entered the canteen, he must have entered the building: $E_1 \rightarrow E_3$
- 3) X must have left the canteen before he left the building: $E_4 \rightarrow E_7$
- 4) Z must have seen X before he left the canteen: $E_9 \rightarrow E_4$
- 5) Y must have sent her e-mail before X could reply to it: $E_8 \rightarrow E_5$
- 6) X must have logged into the network before he could send an email: $E_2 \rightarrow E_5$
- 7) X must have sent an e-mail before he logged out from the network: $E_5 \rightarrow E_6$

Obviously, there is no causal connection between X visiting the canteen and X sending an e-mail.

It may seem that there should be one more connection: X must have entered the building before he could log into the network. This, however, would only be the case if X could not log into the network remotely. This possibility is *not* ruled out in the text. Thus, no causal connection added.

The identified events together with relations between them are shown as a graph in Figure 7 below.




Figure 7. Graph of events

No minimal delays can be identified from the text. Thus, all minimal delays are assumed to be equal 0.

Now, for each event E_i we must find all events that *happened-before* it and all events for which E_i is a *happened-before* event. The results of this activity are given in

Table 2 below.

Table 2. Causal connections between events

<i>Events that happened before</i>	<i>Event</i>	<i>Events that happened after</i>
None	E_1	E_3, E_4, E_9, E_7
None	E_2	E_5, E_6
E_1	E_3	E_4, E_9, E_7
E_1, E_3, E_9	E_4	E_7
E_2, E_8	E_5	E_6
E_2, E_5, E_8	E_6	None
E_1, E_3, E_4, E_9	E_7	None
None	E_8	E_5, E_6
E_1, E_3	E_9	E_4, E_7

Finally, the T_{min} and T_{max} for each event is calculated according to formulae (4). The results are given in Table 3 below.

Note that nothing was deduced about the time of event E_2 . This is because none of the events restricts the time interval for it. This reflects the fact that although X was in the company building, he (or somebody posing as him) could have logged into the company computer network from outside the building. Additional investigation is required to clarify this issue.

Note that the calculated time limits would have been narrower, if the minimal delays between at least some of the events could have been identified. If, for example, we could show that the employee X could not get to the canteen from the building

entrance faster than in d minutes ($E_1 \xrightarrow{d} E_3$), it would have narrowed the time interval for event E_3 by d minutes.

Table 3. Calculated time limits

Event	Description	T_{\min}	T_{\max}
E_1	X entered the building	08:30	08:31
E_2	X logged into the company computer network	$-\infty$	$+\infty$
E_3	X entered the company canteen	08:30	17:31
E_4	X left the company canteen	15:00	17:31
E_5	X sent an e-mail to Y	11:00	$+\infty$
E_6	X logged out of the company computer network	11:00	$+\infty$
E_7	X left the building	17:30	17:31
E_8	Y sent an e-mail to X	11:00	13:00
E_9	Z saw X in the canteen	15:00	17:31

Section 5: Conclusions

A straightforward and useful method for calculating time intervals of events by considering their causal connections with other events whose time is known has been analysed, modelled and discussed. Although the method cannot handle certain complex situations easily, nevertheless, it offers the following major benefits:

- *Simplicity.* The essential ideas of the method can be easily explained to non-scientists including lawyers and jurors.
- *Systematic approach.* The method forces the investigator to analyse causal connections between events more carefully than it would be the case if the reasoning about event times was done intuitively.
- *Can be programmed.* The calculation of time intervals for a given event graph can be performed by a computer program.

Although the method can be used for analysing simple situations like the example given in this paper, it will probably be more beneficial for analysing complex situations with many events and many causal connections. In such situations, the human mind particularly needs a systematic approach to deal with the amount and complexity of the data obtained in an investigation.

Related work

In digital investigations, timestamps can be corroborated in a number of ways. This paper formalised one such way that exploits causal connections between events to establish the order of events and, hence, to restrict possible times of events. Other methods have also been analysed in the literature.

A different approach to using external sources of time is described in (Weil, 2002). It exploits the ability of web servers to insert timestamps into web pages. As a result of this insertion, a web page stored in a web browser's disk cache has two timestamps. The first timestamp is the creation time of the file, which

contains the web page. The second timestamp is the timestamp inserted by the web server. The offset between the two timestamps of the web page reflects the deviation of the local clock from the real time. It is proposed to use that offset to calculate the real time of other timestamps on the local machine. To improve precision, it is proposed to use the average offset calculated for a number of web pages downloaded from different web servers. A general model for this type of reasoning has been defined in Stevens (2004). It links clock settings on different computers through a hierarchical system of clock offsets.

The event time bounding algorithm described in this paper relies on the known causal relationship between events to compute possible time intervals of the events. In complex incidents, however, the determination of causal relationship between events can be difficult due to the large number of possible events. Several techniques, such as Multilinear Event Sequencing (Benner, 1975) and Why Because Analysis (Ladkin, 2000), emerged in the domain of accident investigations to structure that process. The development of similar techniques for the digital forensic domain is ongoing. The reader is referred to Stephenson (2004), Gladyshev & Patel (2004), and Carrier & Spafford (2004) for further discussion.

Finally, it is prudent to note that ideas similar to those described in this paper have been independently discovered by Norwegian researchers and published recently in Willassen, Mjølshes (2005).

Future work

The simplicity of the event model makes the described method straightforward. At the same time, it means that certain situations cannot be handled easily. One such situation arises when an event could have happened in more than one time interval.

Return for a moment to the example from Section 4. Suppose that employee X could have written an email to Y at any time between 11:00 and 17:00, but *not while he was in the canteen*. To represent this fact, we need two time intervals – one before and one after X visited the canteen -- which is not allowed in our model.

One of the ways to deal with this problem is to split the event into two – one event before and one event after X visited the canteen. After time limits calculation, the time intervals for these two events define possible time for the initial (non-split) event. To avoid the need for such a split, it may well be possible to devise a more expressive formalism, which is able to handle multiple possible intervals for a single event. The research of this and other possible improvements of the model presented in this paper is left for future work.

© Copyright 2005 International Journal of Digital Evidence

About the Authors

Dr. Pavel Gladyshev is a College Lecturer in the School of Computer Science and Informatics at University College Dublin (UCD). His research interests are in the area of information systems security and digital forensics. Before joining UCD as a lecturer, Dr. Gladyshev worked as a senior consultant in information systems

security and forensics at the Dublin practice of Ernst & Young. Dr. Gladyshev holds MSc and PhD degrees in Computer Science from UCD and he is a member of the Information Systems Security Association.

Dr. Ahmed Patel received his MSc and PhD degrees in Computer Science from Trinity College Dublin (TCD) in 1978 and 1984 respectively, specializing in the design, implementation and performance analysis of packet switched networks. His research interests spanned topics concerning international networking and application standards, network security, forensic computing, high-speed networks, heterogeneous distributed computer systems and distributed search engines and systems for the Web. Before retirement in spring 2005, he was a Senior Lecturer in the Department of Computer Science, UCD and Head of the Computer Networks and Distributed Systems Research Group.

References

- Bates, J. (1997) Blackmail: case study. *International Journal of Forensic Computing*, 1(2), pp. 9-11.
- Benner, L. (1975) Accident investigation: Multilinear event sequencing methods. *Journal of Safety Research*, 7(2), pp. 67-73.
- Carrier, B.D., Spafford, E.H. (2004) Defining Digital Crime Scene Event Reconstruction. *Journal of Forensic Sciences*, 49(6). Paper ID JFS2004127. Retrieved from <http://www.astm.org/>
- Casey, E. (2004) *Digital Evidence and Computer Crime: Forensic Science, Computers and the Internet*. Second edition, Academic Press.
- Cormen, T., Leiserson, C., & Rivest, R. (1990) *Introduction to Algorithms*, MIT Press.
- Gladyshev, P., Patel, A. (2004) Finite State Machine Approach to Digital Event Reconstruction. *Digital Investigation* journal, 1(2), pp. 130-149.
- Ladkin, B.P. (2000) Causal Reasoning about Aircraft Accidents. *In proceedings of SAFECOMP 2000 conference*, Rotterdam, Netherlands, pp. 344-360.
- Lamport, L. (1978) Time, Clocks, and the Ordering of Events in a Distributed System, *Communications of the ACM*, 7, pp. 558-565
- Stephenson, P., (2004) Application Of Formal Methods To Root Cause Analysis of Digital Incidents, *International Journal of Digital Evidence*, 3(1). Retrieved from <http://www.ijde.org/>
- Stevens, M.W. (2004) "Unification of relative time frames for digital forensics", *Digital Investigation* journal, 1(3), pp. 255-239.
- Weil, M.C. (2002) "Dynamic Time & Date Stamp Analysis", *International Journal of Digital Evidence*, 1(2), Retrieved from <http://www.ijde.org/>
- Willassen, S.Y., Mjøltnes, S.F. (2005) Digital Forensics Research, *Teletronikk* journal, 1, pp. 92-97. Retrieved from <http://www.telenor.com/teletronikk/>

Appendix 1

Claim: If $(A \rightarrow B) \wedge (B \rightarrow C)$ and delays between events A , B and C may be unboundedly small, then the new values of T_{\min}^B and T_{\max}^B will satisfy the definition of T_{\min}^B and T_{\max}^B for arbitrary inputs, only if we use T_{\min}^A in place of T^A and T_{\max}^C in place of T^C in formulae (2).

Proof: First we show that if we use T_{\min}^A in place of T^A and T_{\max}^C in place of T^C in formulae (2), the new values T_{\min}^B and T_{\max}^B will satisfy the definition of T_{\min}^B and T_{\max}^B .

By definition of T_{\min}^A and the temporal property of the *happened-before* relation, $T_{\min}^A \leq T^A < T^B$. By definition of T_{\min}^B , $T_{\min}^B \leq T^B$. Thus, if either T_{\min}^A or T_{\min}^B is selected as the new T_{\min}^B , it will satisfy the definition of T_{\min}^B .

By definition of T_{\max}^C and the temporal property of the *happened-before* relation, $T^B < T^C \leq T_{\max}^C$. By definition of T_{\max}^B , $T^B \leq T_{\max}^B$. Thus, if either T_{\max}^C or T_{\max}^B is selected as the new T_{\max}^B , it will satisfy the definition of T_{\max}^B .

Now we show that if we use values other than T_{\min}^A , and T_{\max}^C in formulae (2), it may lead to incorrect results.

Suppose that we choose some T_{\min}^A in place of T^A other than T_{\min}^A , then it must be true that $T_{\min}^A < T_{\min}^A \leq T_{\max}^A$. If T_{\min}^A is chosen as the new value for T_{\min}^B in formulae (2), then $T_{\min}^A < \text{new } T_{\min}^B$. If the delay between A and B may be unboundedly small, T^A and T^B may be such that $T_{\min}^A \leq T^A < T^B < T_{\min}^B$, which contradicts the definition of T_{\min}^B .

Suppose that we choose some T_{\max}^C in place of T^C other than T_{\max}^C , then it must be true that $T_{\min}^C \leq T_{\min}^C < T_{\max}^C$. If T_{\max}^C is chosen as the new value for T_{\max}^B in formulae (2), then $\text{new } T_{\max}^B < T_{\max}^C$. If the delay between B and C may be unboundedly small, T^B and T^C may be such that $T_{\max}^B < T^B < T^C \leq T_{\max}^C$, which contradicts the definition of T_{\max}^B .

Appendix 2

Given below is the algorithm for steps 5 to 7 of the method described in the paper.

Input data:

- 1) Directed acyclic graph of events $G = (E, H)$, where
 - $E = \{E_j\}$ is a set of events
 - $H = \{h_k\}$ is a set of arcs $h_k = (E_i, E_j)$. An arc between two events exists if and only if $E_i \rightarrow E_j$.
- 2) A minimal delay $d_{E_i E_j}$ is associated with each arc $h_k \in H$. If the delay is unknown, it is assumed equal 0.
- 3) A pair of times $(T_{\min}^{E_j}, T_{\max}^{E_j})$ is associated with each event E_j . For the events whose time is unknown $T_{\min}^{E_j}$ and $T_{\max}^{E_j}$ equal to $-\infty$ and $+\infty$ respectively.

Output data:

A pair of new times $(T'_{\min}^{E_j}, T'_{\max}^{E_j})$ for each event E_j .

Algorithm:

- 1 Compute transitive closure $G^* = (E, H^*)$ of the graph G and the minimal delay for each pair of causally connected events. This can be done, for example, by using a modified Floyd-Warshall algorithm. For a description of the Floyd-Warshall algorithm see (Cormen, Leiserson, & Rivest, 1990).
- 2 For each event E_j in G^* :
 - 3 $T'_{\min}^{E_j} := T_{\min}^{E_j}$
 - 4 $T'_{\max}^{E_j} := T_{\max}^{E_j}$
 - 4.1 For each event E_n in G^* that has an arc connecting it to E_j :
 - 4.1.1 If $T'_{\min}^{E_j} < (T_{\min}^{E_n} + d_{E_n E_j})$ then $T'_{\min}^{E_j} := T_{\min}^{E_n} + d_{E_n E_j}$
 - 4.2 For each event E_k in G^* such that there is an arc from E_j to E_k :
 - 4.2.1 If $T'_{\max}^{E_j} > (T_{\max}^{E_k} - d_{E_j E_k})$ then $T'_{\max}^{E_j} := T_{\max}^{E_k} - d_{E_j E_k}$
- 5 Stop