Persistent instabilities in the high-priority incident workload of CSIRTs

Johannes Wiik
Jose J. Gonzalez
Security and Quality in Organizations research cell
Faculty of engineering and science
University of Agder
4898 Grimstad, Norway
johannes.wiik@gmail.com, jose.j.gonzalez@uia.no

Pål I. Davidsen
Institute for geography
University of Bergen
5020 Bergen, Norway
pal.davidsen@geog.uib.no

Klaus-Peter Kossakowski
Software Engineering Institute Europe
Carnegie-Mellon University
60 322 Frankfurt, Germany
kpk@sei.cmu.org

Abstract

Since their inception Computer Security Incident Response Teams (CSIRTs) have been afflicted by chronic problems concerning workload, quality of service, and sustaining their constituency. We have cooperated with one of the oldest CSIRTs to model the most challenging issues. Low-priority and high-priority incident response cause distinct problems. In a previous paper we dealt with the impact of the exponential growth of low-priority incidents on the CSIRT workload. In this paper we deal with high-priority incident response and its impact on the CSIRT workload and quality of service. One observes long-term instabilities in workload and QoS and, ominously, oscillatory decreasing recognition of the CSIRT by its constituency. An improved communication of the service level provided by the CSIRT is the most effective policy to mitigate long-term instability in the workload and quality of service.

Keywords

Information Security, Incident Response, Incident Management, CERT, CSIRT, Risk Management, System Dynamics

Introduction

This is the second of three papers dealing with chronic problems in a coordinating Computer Security Incidents Response Team (CSIRT). The first paper, entitled "Chronic workload problems in Computer Security Incident Response Teams", and the third paper, entitled "Preserving a balanced Computer Security Incident Response Team constituency", have both been submitted to this conference.

For the sake of brevity we refer to the first paper for most details about CSIRTs in general, the specific coordinating CSIRT and the modeling process, including data mining, model verification, validation and policy testing procedures. In the remainder of this paragraph we only summarize the basic aspects.

The first CSIRT was established in 1988 and quickly new CSIRTs appeared. Since the very beginning CSIRTs have been afflicted by chronic problems. As early as in 1994 a study concluded that the existing CSIRTs were insufficiently funded, understaffed, and overworked (Smith 1994, §3.8.1). A thorough CERT/CC report on the state of the practice of CSIRTs documented that the problems persisted in most of the now approximately two hundred external coordinating CSIRTs over the world (Killcrece et al. 2003). The workload in CSIRTs is overwhelming, implying a wide range of internal problems, such as insufficient funding (van Wyk and Forno 2001; Killcrece et al. 2003), lack of management support, shortage of trained incident handling staff, poorly defined mission and authority, and lack of coordination mechanisms (Killcrece et al. 2003, §3.11). The problems persist – at least, there is no published material that indicates a change for the better. Information available from existing teams indicate no change either.

Among the various types of CSIRTs, coordinating CSIRTs usually have the broadest scope and most diverse constituency among the CSIRT organizational models. A coordinating CSIRT is typically located in a single location, coordinating and facilitating the handling of incidents across a variety of organizations in dispersed locations. We have collaborated with one of the oldest coordinating CSIRTs, and for that matter one of the oldest CSIRTs, to shed light on the causes of the chronic problems and to test management policies. The project was conducted as a PhD thesis and the project team consisted of the PhD candidate, two supervisors from different universities offering a system dynamics programme and the manager of the coordinating CSIRT. The CSIRT staff provided access to the necessary data (mental, written and numerical). For more details of this we refer to the first paper in this series and for an extensive description of the whole project we refer to the PhD thesis (Wiik 2007; Wiik et al. 2009a).

Owing to its restricted availability, the PhD thesis was allowed to disclose the name of the actual CSIRT and to use the real data series. In contrast, this paper and its companion two papers withhold the CSIRT name and use sanitized data – in accordance with the practice in information security.

A crucial finding was that low- and high-priority incidents cause major problems of their own. Low-priority incidents typically are "standard" attacks – such as port scans, spams, fake mails, etc. They are nuisances, more like mosquito swarms, but not serious. Their impact on the CSIRT workload comes from their sheer number and their growth rate, their incidence rate typically doubling in number every year (Killcrece et al. 2003, §3.8.1). The first paper in this series, "Chronic workload problems in Computer Security Incident Response Teams", presented a system dynamics model of low-priority incident and discussed policies to contain the problem. In contrast to low-priority incidents, high-priority incidents are serious but relatively rare. Their impact on the CSIRT performance manifests itself through long-term instabilities in the workload and problems to retain a balanced constituency. This paper deals with the workload instabilities whereas the third paper in this series addresses the constituency aspect.

As for the two companion papers (Wiik et al. 2009a, 2009b), we have modeled the problem in this paper using a new object-oriented system dynamic tool Smia, developed by Dynaplan.

Long-term instabilities in the high-priority incident workload

It is acknowledged that CSIRTs are experiencing large variations in the incident workload (West-Brown et al. 2003; Killcrece et al. 2003). Among the causes of these variations are new vulnerabilities exposing previously secured systems, new malware outbreaks, and new automated tools for exploiting known vulnerabilities, generating waves of new incidents. More long-term fluctuations in workload, in particular those concerning high-priority incidents, have been far less studied and are therefore much less understood.

In contrast to the exponential growth of low-priority incidents, a totally different behavior pattern emerges when one singles out the high-priority incidents. For high-priority incidents reported and sites reporting such incidents the data could imply a long term cyclical pattern after an initial S-shaped growth (Figure 1). While there are incident statistics since 1994, reliable site reporting data are available only for 2000-2005, following the introduction of the pseudo-naming scheme, expanding the previous scheme in place. Despite the lack of site information for 1994-1999, interviews with incidents handlers from this period indicate that the variation in reporting sites had followed the same pattern as the number of high-priority incidents, i.e. S-shaped growth.

S-shaped growth followed by overshoot and oscillations is a generic pattern often found in systems with resources constraints, varying quality of service and long time delays (Sterman 2000, p. 121-122). Indeed, the CSIRT-literature makes a strong case that CSIRTs have resource problems (Smith 1994; van Wyk and Forno 2001; West-Brown et al. 2003; Killcrece et al. 2003). A closer analysis of available information for our CSIRT case indicates that its constituency reacts with long time delays upon variations on the CSIRT's quality of service (Wiik 2007).

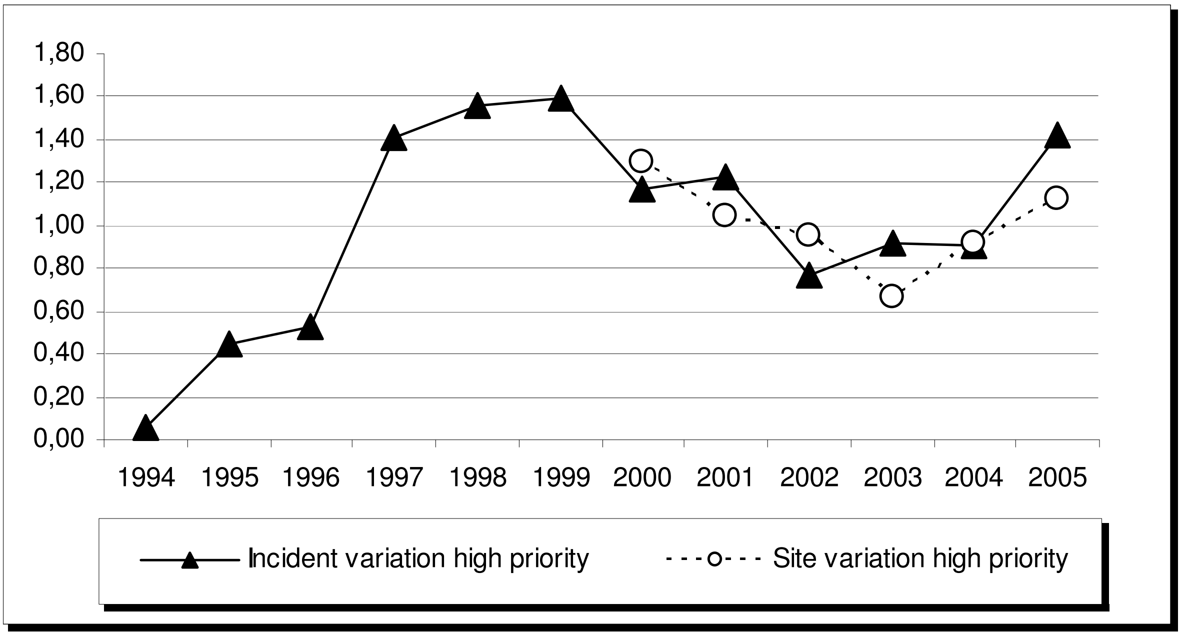

Figure 1 The graph shows the fractional variation of high-priority incidents and reporting sites. To protect the identity of the coordinating CSIRT, the data are shown as fractional variations relative to the average over time

The long-term behavior patterns of incident reporting and reporters presented above raises several research questions:

1. What are the causes for this oscillating behavior in the high-priority workload and in reporting sites?
2. What are the implications of such cycles for incident response effectiveness and the CSIRT mission?
3. How will various policies influence the behavior and the ability to fulfill the CSIRT mission in an effective manner?

The data suggests a relation between the number of sites reporting high-priority incidents and the number of such incidents reported. The reporting frequency measured as incidents reported per site per year was almost constant. That is, the number of reporting sites and the number of incidents are covariant (Figure 1). This suggests that the most important factor driving the high-priority workload is the number of sites reporting.

Notably, a smaller group of sites tended to report at a much higher frequency than the majority that only reported on a very sporadic basis. Accordingly, we distinguish between two groups of reporters: 1) Frequently reporting sites; 2) Sporadically reporting sites.

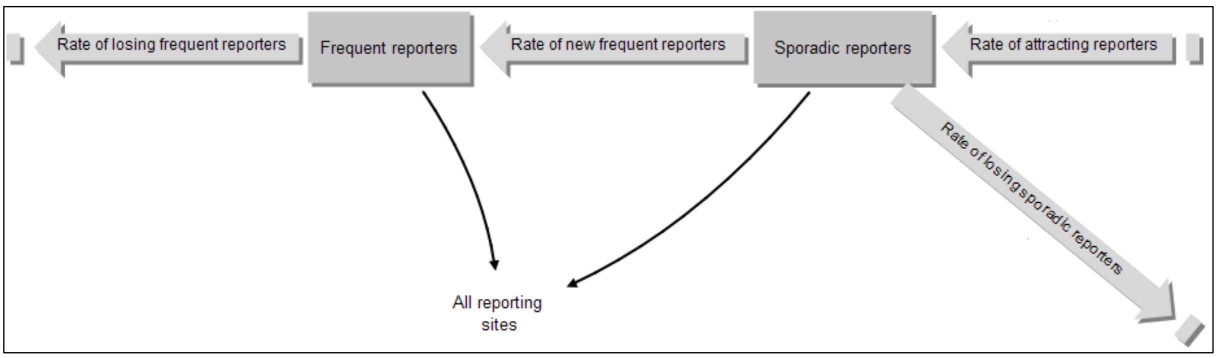

Figure 2 Stock-and-flow structure capturing the number of sites in each category and how they change over time. The square boxes at the beginning or end of flow variables mark the boundary of the model

The number of sites in each category may change because new sites are added to or because sites have been withdrawn from each category. Site categories and how they change are pictured using a stock and flow diagram as shown on Figure 2:

On Figure 3 we have added four feedback loops governing the rate of the four flows.

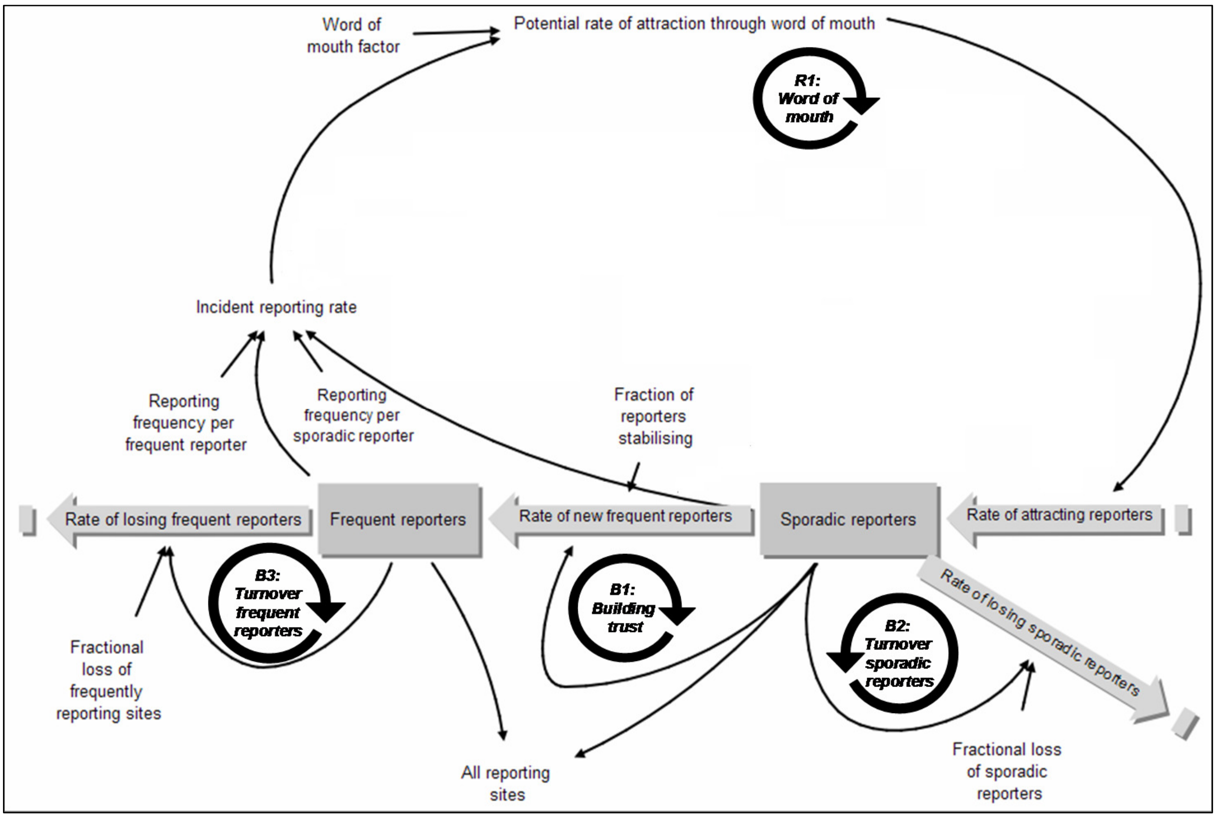

Figure 3 Model structure with focus on reporting sites and word of mouth attracting more sites

The CSIRT is dependent on reporting sites. Sites are attracted to the CSIRT for various reasons at a certain rate and start reporting. As it takes time to build trust between the reporting site and the CSIRT, we assume that all new sites start as what we call sporadic reporters, that is, with a relatively low average reporting frequency.

Sporadic reporters may continue to report on a sporadic basis while a few sites will turn into frequent reporters as they build trust to the CSIRT. Hence, we assume that a certain fraction of sporadic reporters will turn into frequent reporters over time. This process is captured by the variable “Rate of new frequent reporters”, the flow connecting the two stocks, and controlled by the balancing feedback loop “B1: Building trust.” Such a loop will tend to seek an equilibrium point and, if it is the dominating structure, it will exhibit goal-seeking behavior towards this equilibrium.

Frequent reporters are typically other CSIRTs or dedicated security organizations that are able to maintain a steady stream of high-priority incident reports, as opposed to the more ad-hoc reports coming from sporadic reporters. Frequent reporters act typically like funnels, gathering incident information from a wide range of their constituent sites and forwarding the relevant information to other CSIRTs.

The turnover process is governed by the balancing feedback loop “B2: Turnover sporadic reporters.” In a similar manner, there is also turnover among frequent reporters, captured through the balancing feedback loop “B3: Turnover frequent reporters.” Our CSIRT data shows that turnover is generally quite high, with the turnover among frequent reporters being generally lower than for sporadic reporters.

Sporadic and frequent reporting sites report with different frequencies. As mentioned above, the number of incidents reported per time is covariant with the number of sites reporting, and this supports the assumption that the reporting frequency is constant for both kind of reporters. The number of sites in each category and the constant reporting frequencies drive the incident reporting rate to the CSIRT.

Incidents typically involve at least two sites (a source and a victim). More often than not, the source is also a victim being abused to launch an attack against other sites. The reporting sites can be at either end of such an incident, or it can be a CSIRT coordinating the information flow between a site in its own constituency with sites in the constituency of another CSIRT. This means that a wide range of sites may become increasingly aware of each other. The more time the CSIRT has available to handle an incident, the more likely it is that more involved sites will be discovered. Such interactions with various sites can lead to further reports in the future. Sites involved may also recommend the CSIRT's incident response services to others not afflicted by the incident handled. All such activities can lead to what we have referred to as word-of-mouth in the simulation model. I.e., the interaction between the CSIRT and other sites can lead to further attraction of other reporting sites. Hence, incident reporting drives the indicated word-of-mouth and, thereby, the attraction of new sites.

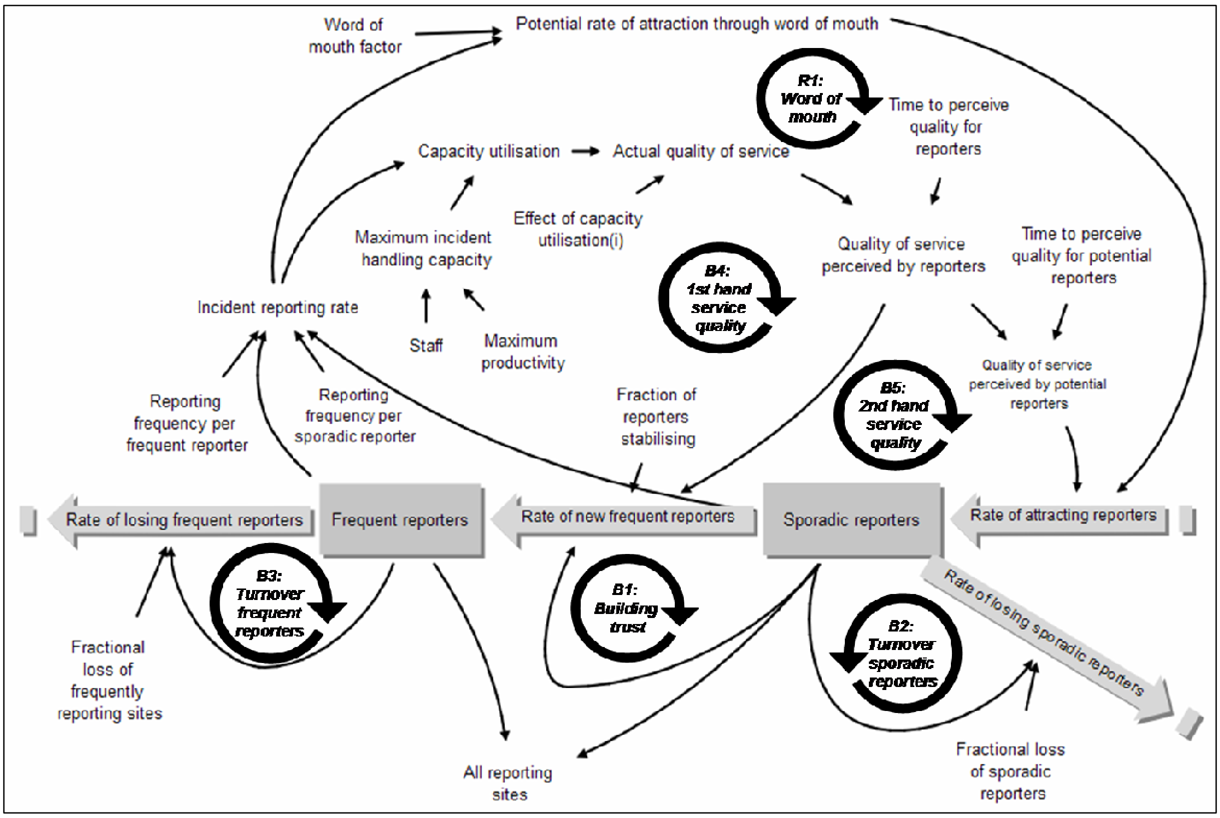

Figure 4 The handling capacity represents an internal limit to the growth of the CSIRTs workload

The word-of-mouth process forms a reinforcing feedback loop, "R1: Word of mouth," where reporting of incidents leads to the attraction of more reporting sites and, thereby, even more reporting. If this loop dominates, it will lead to an exponential increase in the workload. Such growth in reporting is typically experienced by recently established CSIRTs, but it can never go on forever. We assume that the main limit to growth can be found in the CSIRT's internal resources – in practice, its staff. Such a limit plays a crucial role in two balancing feedback loops B4 and B5 concerning perception of service quality by reporters and potential reporters (Figure 4). The maximum incident handling capacity is modeled as a function of maximum productivity and the size of the staff. Both variables are considered constant in this model.

The incident reporting rate and the maximum incident handling capacity determine the capacity utilization, which again influences the indicated quality of service. Quality of service is modeled as an inverse function of capacity utilization. We assume that at maximum capacity utilization the quality of service becomes too low to be useful. Considering prioritization within incident handling such an assumption may sound unreasonable. However, it seems that an overall workload forces incident handlers, despite prioritizing the highest priority incidents, to lower the time spent on all incidents (in order to handle all of them within a reasonable amount of time). (Wiik 2007) Thereby, when the total workload is extremely high, the quality of the handling of almost all incidents will suffer.

Even if a site experiences a low quality of service as a result of high capacity utilization, it does not mean that the perception about quality changes instantly. Rather, the recent and past experience is blended together so that the perceived quality of service gradually changes for those using the service. This means that the perception about service quality is delayed and smoother than actual changes in quality of service. Even smoother and more delayed is the perception of service quality among potential reporting sites that do not have first-hand experience, since their perception is based on reporting sites' perception.

The perceived quality of service for reporting sites influences the rate of new frequent reporters: If sporadic sites do not receive the service they expect, it is less likely that they become frequent reporters. Similarly, the rate of attracting potential reporting sites, 'Rate of attracting reporters', is influenced by the perceived quality of service. Indirectly, this is an effect that is combined with word-of-mouth: Reporting sites simply do not recommend the service as often to potential reporting sites as they would if quality of service had been higher. It is assumed that there is practically no effect from quality of service on the turnover of frequent reporters. The reason for this assumption is that frequent reporters are typically other CSIRTs both outside and within the constituency. Such sites are often less sensitive to low quality of service, since they are experts that need less follow-up per incident. Besides, these sites often have a much better understanding of other CSIRTs' problems. Thereby they recognize the importance of continued reporting despite low quality of service.

The quality of service closes two balancing feedback loops, both inhibiting incident reporting when capacity utilization becomes too high. "B4: First hand service quality" limits the number of sites converted into stable reporters, while "B5: Second hand service quality" limits word-of-mouth and attraction of new sites.

In the appendix we provide the model equations.

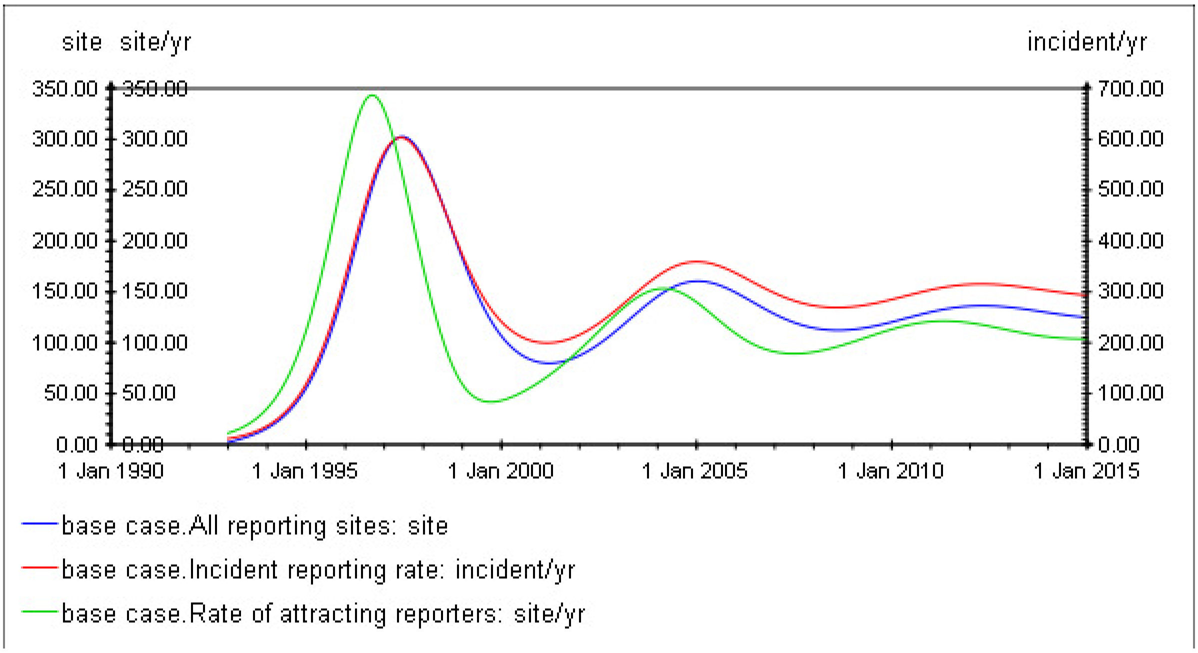

Figure 5 The base case shows a replication of the historical pattern for the incident reporting rate and the number of reporting sites starting with S-shaped growth, with ensuing overshoot and oscillations. The time horizon has been extended to 2015 to assess expected future development based on the model assumptions.

In the base case (Figure 5) we investigate the behavior generated from the structural assumptions above. To get a better picture of the behavior pattern, we have extended the time horizon of the simulation model to 2015, beyond the historical interval (1993-2005) for reporting rates and reporting sites (Figure 1).

From 1993 the number of reporting sites and the reporting rate increase exponentially until around 1997. This growth is generated by the reinforcing feedback loop of word-of-mouth (R1). In the same period, the capacity utilization grows as the workload of the CSIRT increases. This drives down the actual quality of service, but these changes are not recognized immediately by the various sites. With a time lag, the perceived quality of service for sporadic reporters and, next, the perceived quality of service for potential reporters decline at an increasing rate (Figure 6). Interestingly, the rate of attracting reporters peaks around 1997. This point marks a shift after which the reinforcing feedback from word-of-mouth (R1) no longer dominates the behavior. Rather, the balancing loops of service quality (B4 and B5) reduce the attraction of new reporting sites.

From 1997, word-of-mouth (R1) and the gaining of trust from sporadic reporters gradually become weaker, while the various balancing feedback loops controlling reporting site turnover (B2 and B3) gain dominance. This means that the loss of sites gradually starts catching up with the number of sites attracted as well as the number of sites becoming frequent reporters. Hence, the number of sites flattens out and finally peaks around 1999.

From 1999, the turnover of sites is higher than the attraction of new sites and the number of sites reporting starts to decline. Because of the delayed effect from the perceived quality loops (B4 and B5) on the attraction of sites, the perceived quality of service continues to decline, even though the number of sites reporting and thereby the workload is decreasing. This means that there is an overshoot in the workload compared to the available resources created by the delayed response of lower perceived quality of service.

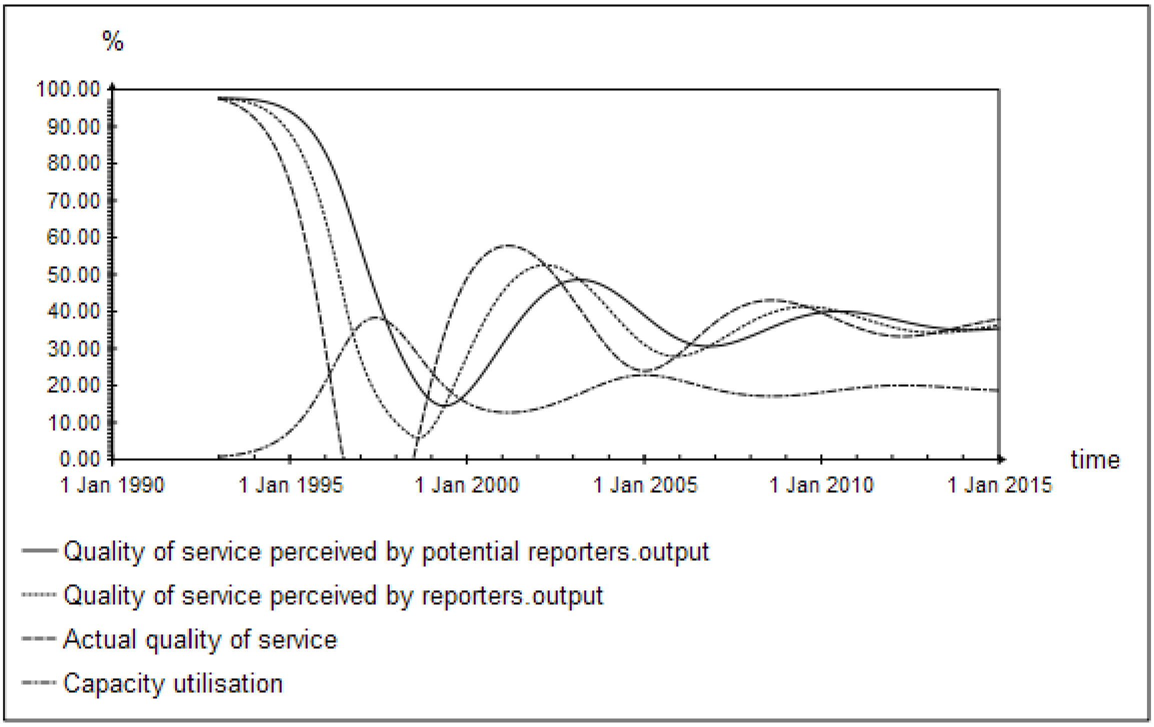

Figure 6 The actual quality of service is driven by the capacity utilization. The quality of service perceived by reporters and potential reporters is smoother and delayed compared to the actual quality of service

The decline in the perceived quality of service and, thus, in the attraction of sporadic sites as well as frequent reporters continues until around 2000. The turnover of sites dominates through B2 and B3. The turnover continues to be higher than the attraction and, consequently, the number of reporting sites continues decreasing until around 2003.

From 2001, a new cycle of growth in reporting sites sets on, driven by word-of-mouth (similar to what happened from 1993). As before, this growth is ultimately counteracted by the decline in perceived service quality and the story repeats itself with damped oscillations for the remainder of the simulation until 2015.

Understanding the structure that causes the fluctuating behavior in the base case scenario provides insights for how to make the incident response service more effective. The CSIRT must strike a balance between the incident workload and the available resources. However, the oscillatory behavior makes the CSIRT continuously undershoot and overshoot such a goal over time. This is a problem for the following reasons:

1. The loss of sites when the workload is too high indicates that the service is not valuable enough to attract new reporting sites, replacing those leaving. As mentioned in p. 7, incidents of lower priority can influence the time spent on higher priority incidents. The total workload will make all sites suffer, and the service quality goes into general decline.
2. If there are excess resources available in incident response, the quality of service will be generally very high. Typically, a relatively new CSIRT with a right skill might provide a very good service, since the capacity utilization is low. The problem is that the unused capacity could have been used to help more sites in need. Furthermore, the situation can easily lead the CSIRT into an overworked situation, driven by word-of-mouth arising from good quality of service.

A CSIRT that is able to maintain a stable equilibrium will be more effective than a CSIRT that experiences long-term fluctuations over- or underutilizing the available capacity. To achieve such a stability goal, we need to identify new policies that will bring the CSIRT closer to such a behavior pattern over time.

Alternative scenarios

We would like to test two scenarios:

1. Adding more resources to avoid overshooting the capacity.
2. Shortening the perception times for acknowledging change in quality of service to adjust faster the workforce and the workload.

The first scenario is implicitly suggested in the literature: CSIRTs are understaffed and underfunded, is a running thread for all major studies of CSIRT performance (Smith 1994; van Wyk and Forno 2001; Killcrece et al. 2003). As to the second scenario, if the oscillatory pattern is created by balancing feedback and long delays, then we need to shorten the delays.

In the first scenario we run the model with 70% additional workforce (the figure 70% is not critical – the point is to show how a marked addition of workforce affects the result). For simplicity, we compare just two variables between the base case and the "Increase staff" scenario: The incident reporting rate and the total number of sites reporting, which are the key variables in the reference behavior graph (Figure 1). Again, we run the model far into the future to evaluate long term impact (Figure 7).

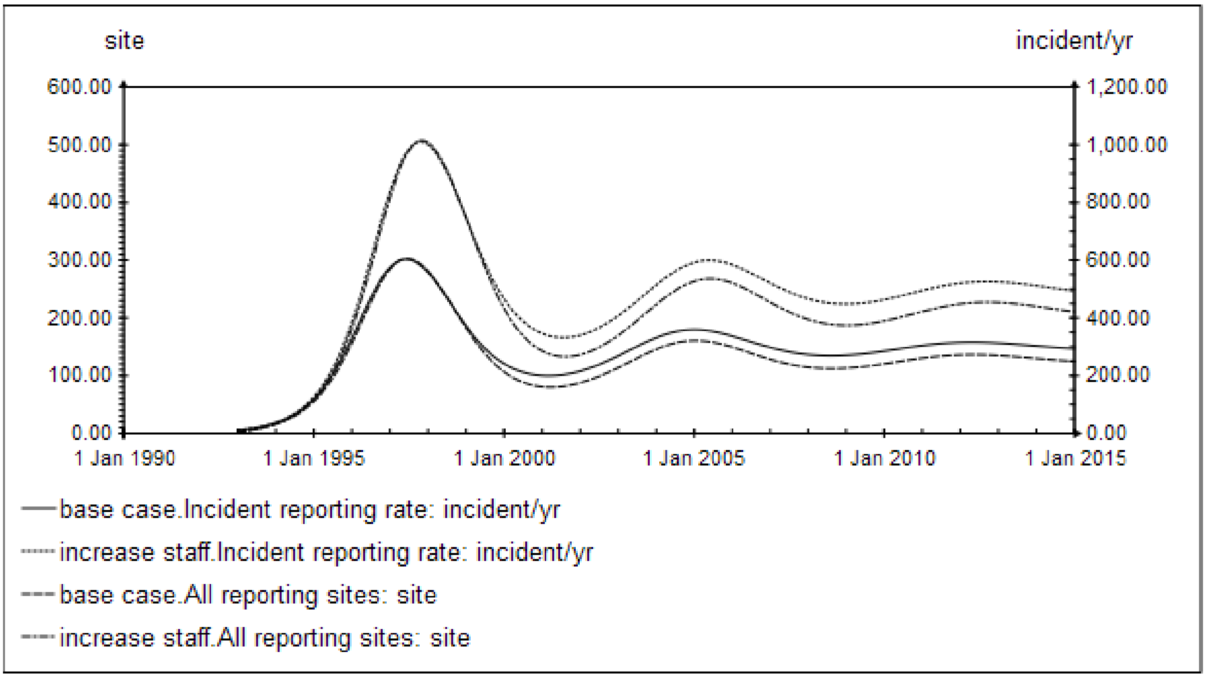

Figure 7 Adding more resources does not solve the problem. Owing to compensating feedback the workload will adjust to the new resources level, and S-shaped growth, overshoot and oscillations will follow with even larger amplitudes

The results show that increasing the size of the staff dedicated to incident response does not change the behavior pattern, simply because the dominating feedback structure remains the same. In fact, the overshoot is even stronger than in the base case. This is because the growth generated from word-of-mouth can build even more momentum before it is controlled by the declining perceived quality of service. The fact that the variations in this scenario are even larger than in the base case, indicate that larger CSIRTs with more staff can experience more problems than smaller CSIRTs. This is because they have larger pools of reporting sites, that generate stronger growth through word of mouth than what might be the case with a smaller CSIRT. The conclusion is, therefore, that adding more resources is not likely to solve the instability problem. Rather, it may make it worse.

In the next scenario we run the simulation after halving the perception times. Expert opinion concurred that that much could, in principle, be accomplished by the CSIRT more actively communicating to potential as well as active reporting sites the service they can expect.

This scenario (Figure 8) shows that, by shortening the delay times, the first overshoot will be much smaller compared to the base case. We also see that the subsequent oscillations are almost completely removed. This scenario indicates that to help as many sites as possible at a satisfactory service level, the long-time delays in the balancing loops "B4: First hand service quality" and "B5: Second hand service quality" should be reduced.

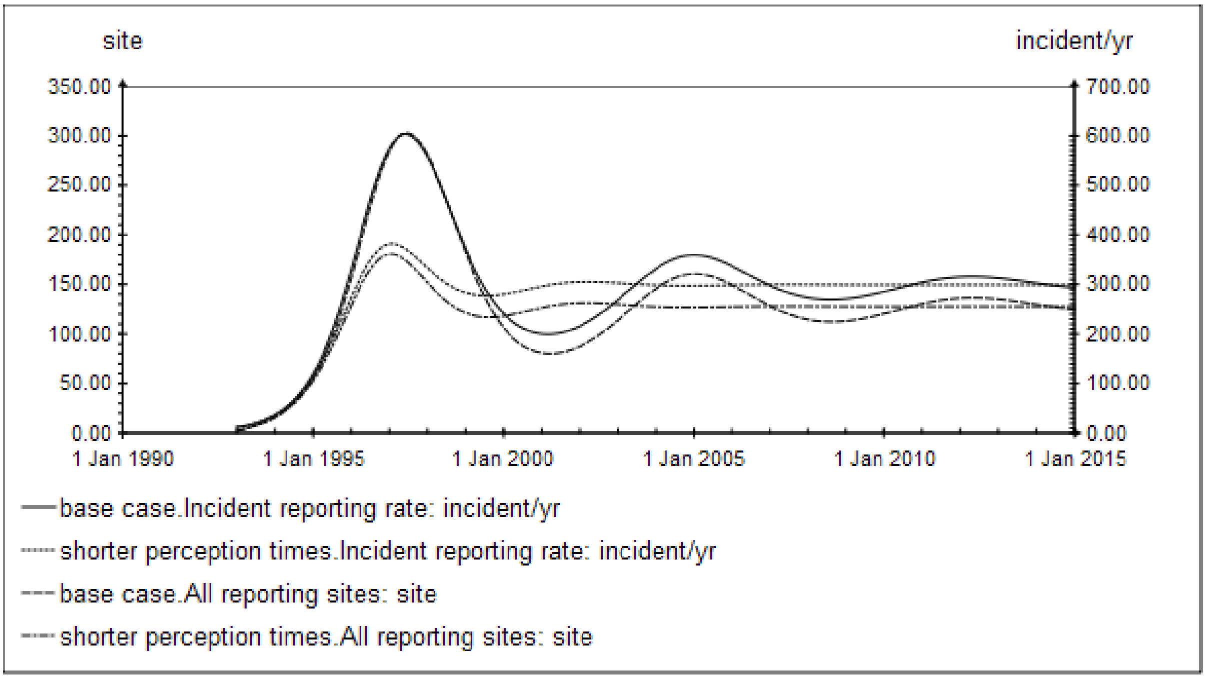

Figure 8 Shortening the perception times tends to stabilize the workload for the CSIRT

Discussion

The model results indicate that CSIRTs can become more effective over time by thinking differently. Some insights were rather expected while others were more surprising.

The discovered behavior pattern for high-priority incidents is surprising. Such behavior has been in the attention shadow, owing to the more conspicuous exponential growth in low-priority incidents that most CSIRTs have experienced. The insight about high-priority incidents comes from the very long time horizon of our CSIRT case. As most CSIRTs have not even existed for that long, it is reasonable to assume that the problem is far from acknowledged by most CSIRTs in operation.

The base-case scenario in the model expresses a hypothesis about long-term instabilities in high-priority incident handling. We have not found any studies of this particular problem in prior research. Hence, we hope that this paper will serve as a contribution for long-term policy analysis.

We do find suggestions in the CSIRT literature that are similar to the strategic goal for achieving effective incident response over time. For example, in the CSIRT Handbook (West-Brown et al. 2003, Section 2.3.3) it is recommended to balance the workload to the available resources by adjusting the range of service offered. To the best of our knowledge, such recommendations are rooted in common sense, rather than in formal theory.

The literature mentions many times that most CSIRTs have been and still are overworked, underfunded and understaffed (Smith 1994; van Wyk and Forno 2001; Killcrece et al. 2003). Our model provides a possible systemic explanation for this phenomenon. As the scenario with added resources to incident response demonstrates, the workload merely adjusts to the new resource level. There is a need for CSIRT services in general that far exceeds any amount of resources that a CSIRT can reasonable expect to have. In other words, adding more resources will only temporarily give the CSIRT a solution until the number of reporting sites and the workload has eaten up any unused capacity.

The key finding is perhaps that reducing the perception delay times will significantly contribute to stabilize the oscillatory behavior of the workload. From an operational perspective, this means that the CSIRT should communicate much more actively what their level of service is. For example, they should announce how effectively they may scan through log files and conduct other time-consuming activities in order to help reporters and other involved sites to mitigate incidents. Thereby, there will be a match between the expectations and the usefulness of the service as seen from the customer point of view. This message can be communicated in any forum where the CSIRT interacts with reporters and users of their service, for example, via the web, e-mail and in workshops.

References

Killcrece, Georgia, Klaus-Peter Kossakowski, Robin Ruefle, and Mark Zajicek. 2003. State of the Practice of Computer Security Incident Response Teams (CSIRTs). Pittsburgh, PA, USA: CMU/SEI.

Smith, Danny. 1994. Forming an Incident Response Team. Paper read at FIRST Annual Conference, at University of Queensland, Brisbane, Australia.

Sterman, John D. 2000. *Business Dynamics : Systems Thinking and Modeling for a Complex World*. Boston: Irwin/McGraw-Hill.

van Wyk, Kenneth R., and Richard Forno. 2001. *Incident Response*. Sebastopol, CA, USA: O'Reilly and Associates Inc.

West-Brown, Moria J., Don Stikvoort, Klaus-Peter Kossakowski, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. 2003. *Handbook of Computer Security Incident Response Teams (CSIRTs)*. 2nd ed. Pittsburgh, PA, USA: CMU/SEI.

Wiik, Johannes. 2007. Dynamics of Incident Response Effectiveness -- A System Dynamics Approach, University of Bergen, Bergen.

Wiik, Johannes, Jose J Gonzalez, Pål I Davidsen, and Klaus-Peter Kossakowski. 2009a. Chronic workload problems in CSIRTs. Paper read at Twenty Seventh International Conference of the System Dynamics Society July, at Albuquerque, NM, USA.

———. 2009b. Preserving a balanced CSIRT constituency. Paper read at Twenty Seventh International Conference of the System Dynamics Society July, at Albuquerque, NM, USA.

Appendix

```
vendor dynaplan
product smia
version 4
language enGB
def {
    submodel model {
        var 'Actual quality of service' = 'lookup linear'('Capacity utilisation','Effect of capacity
            utilisation',false) as %
        var 'All reporting sites' = 'Frequent reporters'+'Sporadic reporters'
        var 'Capacity utilisation' = 'Incident reporting rate'/'Maximum incident handling
            capacity' as %
        var 'Effect of capacity utilisation' = {i=0.0 to 0.3 step 0.3 | 1.0 ,0.0 }
        var 'Fraction of reporters stabilising' = 15 %/yr
        var 'Fractional loss of frequently reporting sites' = 50 %/yr
        var 'Fractional loss of sporadic reporters' = 90 %/yr
        var 'Frequent reporters' = stock 1 site inflow 'Rate of new frequent reporters' outflow 'Rate
            of losing frequent reporters'
        var 'Frequent reporting frequency' = 10 incidents/site/yr
        var 'Incident reporting rate' = ('Frequent reporting frequency'*'Frequent
            reporters'+'Sporadic reporting frequency'*'Sporadic reporters')
        var 'Maximum incident handling capacity' = Staff*'Maximum productivity' as
            incident/yr
        var 'Maximum productivity' = 5 incidents/persons/'work days'
        var 'Potential rate of attraction through word of mouth' = 'Incident reporting rate'*'Word
            of mouth factor'
        component 'Quality of service perceived by potential reporters' = clone #sd.smooth
        component 'Quality of service perceived by potential reporters' {
            var 'delay time' = in:2 'Time to perceive quality for potential reporters'
            var initial = optional in:4{ i == 1 to order | input }
            var input = in:1 'Quality of service perceived by reporters'.output
            var order = optional in:3 1
        }
        component 'Quality of service perceived by reporters' = clone #sd.smooth
        component 'Quality of service perceived by reporters' {
            var 'delay time' = in:2 'Time to perceive quality for reporters'
            var initial = optional in:4{ i == 1 to order | input }
            var input = in:1 'Actual quality of service'
            var order = optional in:3 1
        }
        var 'Rate of attracting reporters' = flow 'Potential rate of attraction through word of
            mouth'*'Quality of service perceived by potential reporters'.output
        var 'Rate of losing frequent reporters' = flow 'Frequent reporters'*'Fractional loss of
            frequently reporting sites'
        var 'Rate of losing sporadic reporters' = flow 'Fractional loss of sporadic
            reporters'*'Sporadic reporters'
        var 'Rate of new frequent reporters' = flow 'Sporadic reporters'*'Fraction of reporters
            stabilising' * 'Quality of service perceived by reporters'.output
```

```
var 'Sporadic reporters' = stock 1 site outflow 'Rate of new frequent reporters' outflow 'Rate
    of losing sporadic reporters' inflow 'Rate of attracting reporters'
var 'Sporadic reporting frequency' = 1.5 incident/site/yr
var 'Time to perceive quality for potential reporters' = 12 mth
var 'Time to perceive quality for reporters' = 12 mth
var 'Word of mouth factor' = 1site/incident
unit 'incident/person/day' = incident/persons/dy
unit 'incident/site/year' = incident/site/yr
unit 'site/incident' = site/incident
unit 'work day'&'work days' = unit (1/210) yr
var Staff = 1.5 persons
submodel globals {
    var 'exchange rate'&exchrate = {r ≤ text|"EUR" =>1}
    var 'game step' = 'time step'
    var 'report step' = 1 yr
    type 'role list' = [ ]
    var 'start future' = 'start time'
    type horizon = date(1993,Jan) to date(2015,Jan) step 1 mth
    submodel sd {
        component 'pipeline delay vector'&delaypplv {
            var 'delay time' = in:2 1*'time step'
            var initial = optional in:3 input
            var input = in:1 flow { 0.0, 0.0}
        }
        component 'pipeline delay'&delayppl {
            var 'delay time' = in:2 1*'time step'
            var initial = optional in:3 input
            var input = in:1 flow 0.0
        }
        component 'sample if' {
            var condition = in:1 false
            var initial = optional in:3 input
            var input = in:2 linear flow 0.0
        }
        component 'sliding average' {
            var initial = optional in:3 (input-input)*period
            var input = in:1 0.0
            var period = in:2 'time step'
        }
        component 'sliding integrate' {
            var initial = optional in:3 (input-input)*period
            var input = in:1 0.0
            var period = in:2 'time step'
        }
        component Euler {
            var input = in:1 0.0
        }
        component delay {
            var 'delay time' = in:2 'time step'
            var initial = optional in:4 { i == 1 to order | input*'delay time each' }
```

```
                var input = in:1 flow 0.0
                var order = optional in:3 1
            }
            component derivate {
                var input = in:1 0.0
            }
            component init {
                var input = in:1 0.0
            }
            component integrate {
                var initial = optional in:2 (input-input)*'time step'
                var input = in:1 0.0
            }
            component sample {
                var first = in:2 'start time'
                var initial = optional in:4 input
                var input = in:1 0.0
                var interval = in:3 'time step'
            }
            component smooth {
                var 'delay time' = in:2 'time step'
                var initial = optional in:4{ i == 1 to order | input }
                var input = in:1 0.0
                var order = optional in:3 1
            }
        }
    }
  }
}
```