Rogue Femtocell Owners: How Mallory Can Monitor My Devices

David Malone, Darren F. Kavanagh and Niall R. Murphy

***Abstract*—Femtocells are small cellular telecommunication base stations that provide improved cellular coverage. These devices provide important improvements in coverage, battery life and throughput, they also present security challenges. We identify a problem which has not been identified in previous studies of femtocell security: rogue owners of femtocells can secretly monitor third-party mobile devices by using the femtocell's access control features. We present traffic analysis of real femtocell traces and demonstrate the ability to monitor mobile devices through classification of the femtocell's encrypted backhaul traffic. We also consider the femtocell's power usage and status LEDs as other side channels that provide information on the femtocell's operation. We conclude by presenting suitable solutions to overcome this problem.**

***Index Terms*—Femtocell, security, traffic analysis, cellular devices, rogue owners.**

I. Introduction

Femtocells allow operators to provide improved coverage in a cellular network by providing low-power base stations which can be installed in homes and offices. The standardisation and security of femtocells has received considerable attention recently [1], [2]. The femtocell uses a customer's broadband connection to backhaul traffic to the mobile operator's network. Accordingly, this creates security concerns, and consequently femtocells use security protocols such as IPsec [3], which encrypt IP traffic that is sent to or from the Mobile Network Operator (MNO). Customers may also be concerned about a third party's device connecting to their femtocell, because, for example, the customer typically pays for the backhauled traffic. Consequently, femtocells have an access list (ACL) feature whereby only cellular devices with phone numbers from a configurable set can connect. This is sometimes referred to as a *closed access* femtocell [4].

In this paper, we consider the prospect that a malicious customer, Mallory, uses their femtocell to monitor mobile devices, belonging to a third party, Alice. By adding a device to Mallory's femtocell's ACL, Mallory can secretly allow Alice's device to backhaul traffic through Mallory's network without Alice's knowledge. While this traffic will be encrypted and authenticated, IPsec does not provide protection against traffic analysis, and so Mallory will be able to make certain inferences about the mobile device. For example, Alice might be Mallory's neighbour, and by adding Alice's phone to the ACL for Mallory's femtocell, Mallory can attempt to identify when Alice makes calls, by observing patterns in the encrypted traffic. In such an attack, the femtocell now acts like a *rogue femtocell*, because Mallory controls the network used for backhaul, and therefore can capture the traffic, albeit in encrypted format. This attack can occur without Alice's knowledge or consent. In contrast to other recent attacks on femtocells [5], this requires no modification of the femtocell's hardware or software.

D. Malone is with the Hamilton Institute at National University of Ireland, Maynooth. D.F. Kavanagh is with the University of Oxford. N.R. Murphy is a Site Reliability Engineer at Google. This work is supported by Science Foundation Ireland under Grant No. 08/SRC/I1403 and 07/SK/I1216a.

Femtocell security was considered during its standardisation [1], [6], however issues still remain [2], [4], [7]. In [8] the authors classify a number of attacks on femtocells. The attack we describe here has similarities to attack 5 (man-in-the-middle) listed in [6], [8]. However, our attack differs in a significant way, because it is entirely passive and not foiled by authentication of the femtocell or the application of cryptography to communication between the femtocell and the MNO's network. While traffic analysis has been considered on the *air interface*, traffic analysis of the *backhaul network* is easier to achieve, and requires no specialist equipment. A computer and an Ethernet hub, or similar device, are sufficient to perform traffic analysis. From the literature, it appears that this new attack has not been studied before or included in previous classifications of attacks on femtocells. As this attack was not identified during standardisation or documented in the literature, it seems likely that femtocell deployments are not being secured against this attack.

Fig. 1 illustrates the system architecture, showing how the femtocell that is to be subject to traffic analysis integrates into the cellular network. Again, note that this set-up does not require any attack on the femtocell device (hardware) itself, instead packets *to* and *from* the femtocell are passively monitored. This traffic is labelled as DST (destination) and SRC (source), respectively. Likewise, monitoring of the femtocell's power usage or status LEDs is passive and requires no modification of the femtocell.

In the following sections, we demonstrate this attack using traffic traces from a live femtocell. We identify basic traffic characteristics that can be used in the traffic analysis. Using these, we derive features that can identify common activities of mobile devices, such as sending SMS (short message service) messages, making calls, web activity, etc. We then demonstrate that by using multiple femtocells we can correlate calls and so potentially identify who is involved in the calls. We also consider power usage and the femtocell's status LEDs as sources of additional information. Finally, we propose solutions to mitigate this problem.

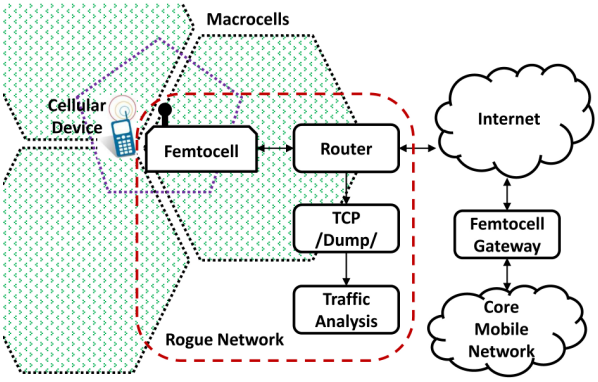

Fig. 1. Rogue femtocell setup, showing integration into the mobile network.

II. Analysis of Single Femtocell Traffic

In this section, we first describe our basic experimental setup. Then we look at the features available to us for traffic analysis of the backhauled traffic from a single femtocell, and show how these relate to the activity of a device on the femtocell. Finally, we show a simple scheme for traffic classification based on these features.

A. Experimental Setup

The femtocell is an Alcatel-Lucent 9361 Home Cell V2-V. It has a 3G radio interface and connects to a customer's network using standard Ethernet. It was connected to a residential router running FreeBSD, and traffic to/from the femtocell was collected using the *tcpdump* tool. This femtocell uses IPsec-over-UDP to communicate with the mobile operator's network. It also generates a small amount of non-encrypted traffic, including NTP time-synchronisation packets and traffic similar to a traceroute. The analysis herein focuses on encrypted traffic. When the femtocell is idle, there are 1–3 packets per second (PPS) of encrypted traffic. Using the femtocell's ACL, we restrict its use to a single device, a Nokia X6 smartphone. *Test Procedure:* To demonstrate the attack, we generate five different types of traffic using a phone connected to the femtocell. Four types of traffic are generated by using the phone for SMS messages, MMS messages, phone calls and web browsing. The last type is generated by taking actions that relate to the network (e.g. moving the device outside the coverage range of the femtocell, turning the phone on/off, ...). We recorded timestamps for the start and end of the activities. These timestamps represent ground-truth for the events up to some small gap between, say, pressing "off" and the phone powering down. We describe our analysis of the packet traces, in comparison to the logged events.

B. Traffic Features

We consider three sources of information for the traffic analysis. The first is traffic timing information, as we record accurate timestamps for the arrival of each packet. We expect that, for example, during a voice call we will see frequent packet transmissions. The second source of information we consider is packet size. We expect to see small packets for a voice call or an SMS message, while web-browsing or multimedia messages (MMS) tend to generate large packets. The final source of information that we consider is the IP Type Of Service (TOS) field (or Differentiated Services Code Point). This field indicates if the packet wants special treatment from the network (e.g. Expedited Forwarding), and so cannot be encrypted. We expect that voice packets may request higher priority handling than, say, an SMS message.

Figure 2(a) shows the number of packets per second received/sent by the femtocell over a traffic trace. The number is calculated using bins covering a duration of 5s. The recorded timestamps for various events, such as the start/end of calls, start/end of web browsing, sending of messages, etc. are shown using stem plots below the axis. The plots show a good correspondence between the activities and the number of packets sent. As we expect, events such as web browsing or calls correspond to high PPS rates, and other events cause small but noticeable increases above the baseline. Note that not all the network management events, such as joining/leaving the cell, cause a change in the number of packets per second. In some cases a delay is observed: powering on a phone may result in network activity several seconds later. This may be because the phone does not immediately connect to the femtocell after booting.

Figure 2(b) shows the size of packets sent over time. A clear data signature is present here, showing that during web browsing a significant number of large packets are sent and received. Similarly, as we expect, there is a pattern of large packets being sent *but not received* when an MMS is transmitted. There is also a clear pattern in the packet sizes being sent at the start of each call, during each call and when an SMS is sent. This is highlighted in Fig. 3.

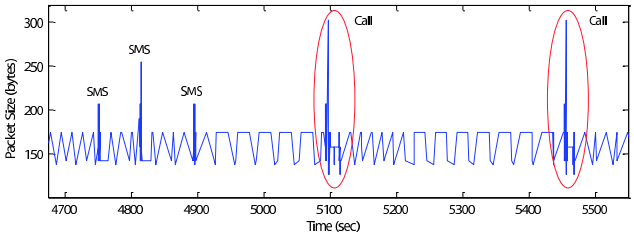

Fig. 3. Close-up of packet sizes that are produced during a phone call.

Moving to Figure 2(c), which shows the frequency of particular TOS values over 10s intervals. Visual inspection of these feature shows strong correlation between the different data classes and TOS activity. For example, the *TOS SRC 72* feature shows strong activity for the start and end of phone calls, *TOS SRC 184* shows strong activity during calls and *TOS DST 0* shows strong activity for phone calls and web browsing. Other values don't show an obvious correlation with the events, e.g. *TOS SRC 192* has a periodic nature, possibly indicating that this is used for network monitoring or reporting.

C. Automatic Classification

Based on our findings, we devised a threshold-based classification algorithm operating on 10s intervals, without keeping

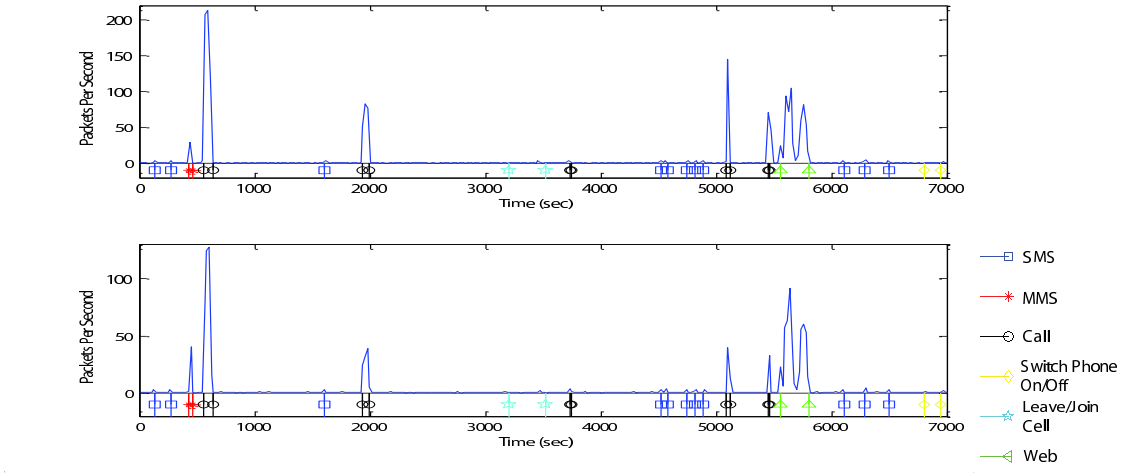

(a) Packets per second vs. time for traffic to the femtocell (above) and from the femtocell (below), calculated with 5s bins.

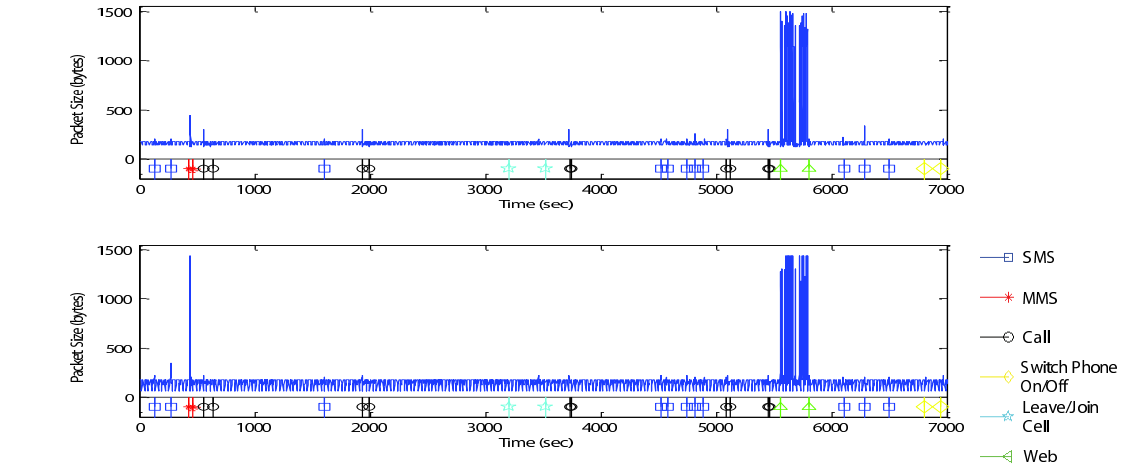

(b) Packet size vs. time for packet to the femtocell (above) and from the femtocell (below).

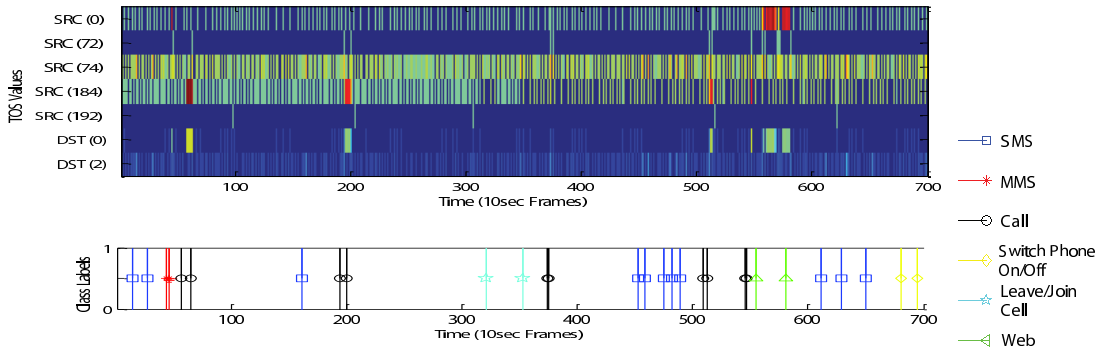

(c) TOS field labels frequency vs. time for packet to/from the femtocell. 10s bins.

Fig. 2. Features identified for Traffic Analysis. Bin sizes are selected to provide clear features.

state between intervals. Threshold values were inferred from empirical observations of real data traces. Pseudocode for this algorithm is shown in Figure 4. The algorithm begins by discarding certain background traffic, based on packet size and TOS. It then looks for signatures for each activity based on numbers of packets with particular size and TOS.

Comparison of the output of the algorithm and the logged list of events indicates that even this simple algorithm can

```
for each (10s interval) {
  Remove background traffic (size, TOS, direction)
  Count number_of packets for each (TOS, direction)
  Store largest packet size for each (TOS, direction)
  if (number_of (TOS 184,SRC) packets > 1)
    event "Call in progress";
  if (number_of (TOS 0,SRC) packets > 0) {
    if (largest (TOS 0, SRC|DST) > 800)
      event "Web session in progress";
    else if (largest (TOS 0, DST) > 800)
      event "Recv MMS in progress";
    else if (largest (TOS 0, SRC) > 800)
      event "Send MMS in progress";
    else
      event "Small Data/MMS in progress";
  }
   if (number_of (TOS 74) > 0 &&
       number_of (TOS 0|72|184, SRC) == 0)
      event "Signaling or SMS";
}
```

Fig. 4. Stateless algorithm for classification of activity.

identify when events occur and can identify events with good accuracy, despite some ambiguity between Signaling/SMS, Web/MMS and pre-call signaling/SMS. The ambiguity around Web/MMS and signaling around a call can be resolved easily by looking at a series of 10s intervals. Resolving the signaling/SMS ambiguity may be more challenging, as SMS messages are a specialised form of signaling in some networks. A small number of events are not identified, however these events appear to generate little or no traffic (e.g. leaving the femtocell's coverage area). With these provisos, over 15000s of trace we can correctly identify over 35 events with only one false positive.

We have established that using features such as *PPS*, *Packet Size* and *TOS values* collectively, it is possible to classify the majority of events. While specific details of these might change in a different femtocell implementation, our analysis suggests that the attack will remain feasible. It is not difficult to postulate a supervised or unsupervised classification algorithm to automatically classify the cellular activities.

III. Two Femtocell Traffic Analysis

In this section, we consider traffic analysis, where the attacker has access to two femtocells and monitors both, to observe temporal correlations of traffic. Though more challenging for an attacker, this could be used to to identify when pairs of target users are in communication.

Our experimental setup is similar to that in Section II, except that we now monitor two femtocells, which we label Femto 1 and Femto 2, for a period of approximately one hour. Monitoring takes place at two distinct routers, whose clocks are synchronised by NTP. In this experiment, the phones in use at both ends are Apple iPhone 3Gs. A series of SMS/Calls/MMSes are made between a user at Femtocell 1 to either a user at Femtocell 2, or users external to these femtocells. A list of events is recorded at Femto 1, with approximate timestamps shown in Table I.

When we remove the background traffic, the time history of packet sizes sent from Femto 2 or packets sent to Femto 1 is shown in Figure 5. We plot the packets at Femto 1 above the axis, and those at Femto 2 below the axis, so a comparison can be drawn without the points being overlaid. Though not shown here, time histories for packets from Femto 1 and to Femto 2 show similar features.

Time	Event
0	SMS with internal and External
160	SMS with between Femto 1 & 2
200	Web browsing at Femto 1
270	Call from Femto 1 to External
390	Call from Femto 1 to Femto 2
540	SMS from Femto 1 to External
700	SMS to Femto 1 from External
1060	SMS from Femto 1 to External
1100	Send MMS and SMS from Femto 1 to Femto 2
1280	SMS from Femto 2 to Femto 1
2690	Call from Femto 1 to Femto 2
2820	SMS to Femto 1 from External
3280	Call from Femto 1 to External

TABLE I
Events recorded at Femto 1, with approximate timestamps.

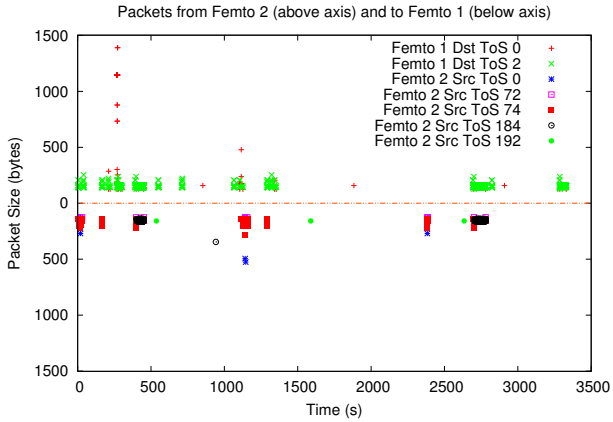

Fig. 5. Time history of packet sizes and TOS from Femto 2 and to Femto 1. The sizes of packets observed at Femto 1 are shown above the axis and Femto 2 are shown below the axis.

If we look at the events listed in Table I, we see that for events just involving Femto 1, there is only traffic above the axis. For events involving both Femto 1 and 2, there is similar traffic observed at both Femtocells. For example, at around 270s we see a number of large packets at Femto 1, but no traffic at Femto 2. This corresponds to a call at Femto 1 to an external number. Similarly, just before 500s, we see a period with a long stream of QoS-marked packets at both Femto 1 and Femto 2. This corresponds to a call through both femtocells.

We use the classifier from Figure 4 to identify the activity in both Femtocells and then align the classified intervals. We find that the classifier identifies all events recorded at Femto 1, plus a small number of additional events at Femto 2 (for example, a short data session at about 2380s). The events which involve both Femto 1 and Femto 2 are identified at both ends, and show strong temporal correlation.

While this demonstrates that it is easy to identify calls and other events that are common to both femtocells, over a long period, it is easy to imagine that there will be some false positives. However, using finer details of each event, it may

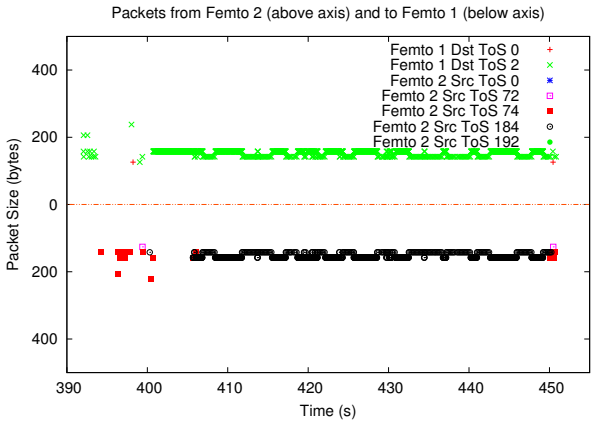

Fig. 6. Time history of packet sizes and TOS from Femto 2 and to Femto 1 during a call. The sizes of packets observed at Femto 1 are shown above the axis and Femto 2 are shown below the axis.

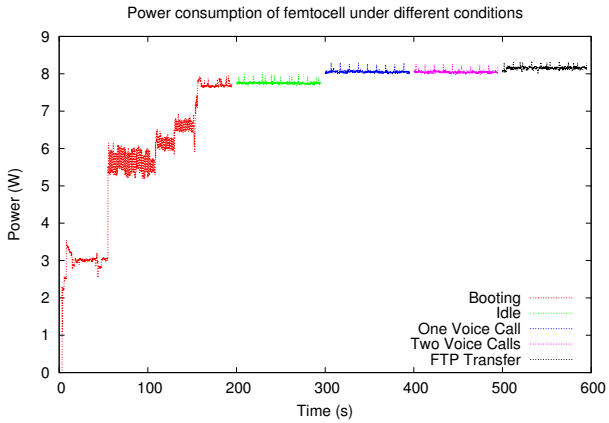

Fig. 7. Power usage of the femtocell under different conditions.

be possible to eliminate these false positives. For example, Figure 6 shows the time history of the call which ends just before 500s. We can see that once the call is fully established around 405s, there is a strong correlation between the packet sizes sent from Femto 2 and received at Femto 1. This pattern is unlikely to be present if we had matched two unrelated calls which happened to be contemporaneous. This allows us to identify calls between pairs of users with high confidence.

IV. Other Side Channels

Beyond traffic analysis, other side channels can be used to make inferences about the femtocell, which do not require any modifications. While traffic analysis provides rich information about activity on the femtocell, other side channels could be used in the absence of traffic information, or to complement traffic analysis. The two sources of side information we consider here are power usage and the status LED.

A. Power consumption

In this section we consider a single femtocell, as described in Section II. We now monitor power usage at the DC input on the femtocell using an Energino [9]. The Energino is an Arduino-based device developed by CREATE-NET which provides fine-grained high-resolution energy usage measurements. It powers a device over a standard DC jack and requires no modification of the monitored device.

Figure 7 shows time histories of power usage of the femtocell when (1) the femtocell is booting; (2) the femto cell is idle; (3) one call is active; (4) two calls are active and (5) a data transfer is in progress. Note, gaps are shown between each time history as the five time histories were not contiguous.

While power usage in situation 3 (one call) and 4 (two calls) are similar, otherwise there is a significant difference in mean power usage of at least 0.1W (at 95% confidence). As the variance of the power usage is small after boot time, it would be possible to use hypothesis testing to classify the activity on the device. A more complete model of femtocell power usage is built in [10]. We conclude that the presence and absence of calls could be determined using power measurements alone.

Fig. 8. Status LEDs. Above, LEDs show power and connection to the MNO are available. Below, device is using the femtocell for call or data session.

B. Status LEDs

A number of status LEDs are present on the femtocell, including two Ethernet LEDs, a power LED, a *system* LED and a *phone* LED. Figure 8 shows the latter three LEDs. The documentation indicates that the system LED is lit when a connection has been established to the MNO. This seems to indicate that the IPsec session has been established and higher-layer communication is in progress.

Again, the documentation indicates that the phone LED is lit when a phone is in use. In practice, it seems to be lit whenever a device using the femtocell has an active voice/data session with the femtocell. When used in combination with the femtocell's ACL, this LED allows us to determine with certainty if a device on the access list is in range. Though limited, this information could be combined with traffic analysis to eliminate false identification of events.

V. Discussion

While we have focused here on identifying types of activity, we note that there have been recent advances in attacks on encrypted voice over IP. For example, in certain circumstances the size of packets from a codec that compresses voice can be used to identify spoken phrases [11]. Such an attack could be applied in a femtocell setup. Another extension is to go

beyond a passive traffic analysis attack by actively generating traffic, say by calling the monitored device, to determine if it is in range of the femtocell. The monitoring described here could also be a prelude to an active attack.

We have seen that monitoring the temporal correlation of traffic at multiple femtocells could also be used to identify pairs of target users that are in communication. Such relatively minor extensions of our basic attack highlight its potential seriousness with respect to users' privacy. For example, consider two celebrities, whose phone numbers may be known by the press. By setting up femtocells near each celebrity's residence, an attacker can determine when calls and SMS exchanges take place between them, resulting in a significant loss of privacy.

These problems arise because although the femtocell's owner has authorised a device to use the femtocell, the device's owner has not authorised the device to connect via the femtocell. In WiFi networks, this problem is known as the rogue access point problem [12]. When we informally surveyed a number of femtocell offerings, we found no evidence of systems that require authorisation from devices' users. Note that the ACL feature actually exacerbates the problem, because it allows targeting of devices. By restricting access to the femtocell to targetted devices, the attacker can eliminate traffic from other, uninteresting, cellular devices in the vicinity. This means that we only have to consider one device per-femtocell, which simplifies the use of traffic analysis or side channels.

Given the seriousness of this attack, it seems important to provide protection against such traffic analysis. Since seamless integration is important for femtocell deployment, any solution must be either transparent to the users or simple to use, as *zero touch installation* is target feature of femtocells [4]. We propose three approaches to mitigate such attacks.

Mitigation Option "Dummy Traffic": One option is for the femtocell to continuously generate traffic. However, generating dummy traffic that does not leak any information about a user's activity could be challenging.

Mitigation Option "IMEI/IMSI Verification": An alternative is to make it harder to add a device to the femtocell's ACL without having legitimate access to the device. A basic ACL just requires knowledge of the devices cell phone number. Requiring additional details, such as the International Mobile Equipment Identity (IMEI) or International Mobile Subscriber Identity (IMSI) number, increases the barrier to maliciously adding a device. The IMEI number, which identifies the device, is easily accessible on most devices and MNOs, and users are familiar with the idea of using it for blacklisting stolen phones, for example. However, if access to the femtocell is tied both to the phone number and the IMEI, the user will have to update the femto's ACL whenever they change their device. The IMSI, which identifies the device's user via the SIM card, seems like a better match. However, in practice the IMSI number is not as broadly accessible on devices and is not commonly used by end users.

Mitigation Option "User Verification": This option explicitly checks that the device's user wants to use the femtocell. When a user is added to a femtocell's ACL, an SMS message could be sent to the user saying "You have been added to Mallory's femtocell. If you trust Mallory and want to use the femtocell, reply to this text." In general, users are already familiar with such explicit opt-in texts. Subsequent management of the list of opt-in femtocells could be provided via a web interface.

These mitigation techniques provide protection against our passive attacks, though they are not effective against an attacker who modifies the femtocell's hardware or software [2], [5], [13]. The *Dummy Traffic* option is unlikely to be a practical solution to the passive attacks, as it is wasteful of backhaul resources and would be less likely to be accepted by consumers who are paying for data backhaul. Using *IMEI/IMSI Verification* or explicit *User Verification* both require the user to interact with their phone, and might be subject to to social engineering attacks or circumvented if the attacker has temporary access to the device. Overall, we believe *User Verification* would offer a good trade-off between usability and security for most users and addresses the fundamental issue of the user consenting to use the femtocell.

VI. Conclusion

We have shown that femtocells can pose a significant privacy risk through monitoring of targetted third party cellular devices. This shortcoming has been overlooked in the literature. Using empirical data analysis of real femtocell traffic, we show how to classify between different user activities. We also consider power usage and status LEDs as an additional source of information. Lastly, we have proposed solutions to mitigate against this security problem.

Acknowledgments — We thank Roberto Riggio at CREATE-NET for the power measurements in our femtocell testbed and Lesley Malone for assistance with the two femtocell tests.

References

[1] D. Knisely, T. Yoshizawa, and F. Favichia, "Standardization of femtocells in 3GPP," *IEEE Communications Magazine*, vol. 47, no. 9, pp. 68–75, 2009.

[2] S. Gold, "Cracking cellular networks via femtocells," *Network Security*, vol. 2011, no. 9, pp. 5–8, 2011.

[3] S. Kent and K. Seo, "Security architecture for the internet protocol," RFC 4301, 2005.

[4] L. Mohjazi, M. Al-Qutayri, H. Barada, K. Poon, and R. Shubair, "Deployment challenges of femtocells in future indoor wireless networks," in *IEEE GCC Conference and Exhibition (GCC11)*, pp. 405–408.

[5] R. Borgaonkar, N. Golde, and K. Redon, "Femtocells: a poisonous needle in the operator's hay stack," *Black Hat Las Vegas*, 2011.

[6] 3GPP, "Technical specification group services and system aspects; security of H(e)NB," http://www.3gpp.org/ftp/Specs/archive/33_series/33.820/33820-830.zip, Tech. Rep. 3GPP TR 33.820.

[7] I. Bilogrevic, M. Jadliwala, and J. Hubaux, "Security issues in next generation mobile networks: LTE and femtocells," in *2nd International Femtocell Workshop*, 2010.

[8] C.-K. Han, H.-K. Choi, and I.-H. Kim, "Building femtocell more secure with improved proxy signature," in *IEEE GLOBECOM*, 2009, pp. 1–6.

[9] K. Gomez, R. Riggio, T. Rasheed., D. Miorandi, and F. Granelli, "Energino: A hardware and software solution for energy consumption monitoring," in *International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2012, pp. 311–317.

[10] R. Riggio. and D. Leith, "A measurement-based model of energy consumption in femtocells," in *Wireless Days 2012*, 2012.

[11] C. Wright, L. Ballard, S. Coull, F. Monrose, and G. Masson, "Uncovering spoken phrases in encrypted voice over IP conversations," *ACM Trans. Info. and System Security*, vol. 13, no. 4, pp. 1–30, Dec 2010.

[12] K. Hole, E. Dyrnes, and P. Thorsheim, "Securing Wi-Fi networks," *IEEE Computer*, vol. 38, no. 7, pp. 28–34, 2005.

[13] R. Borgaonkar, K. Redon, and J. Seifert, "Security analysis of a femtocell device," in *International Conference on Security of Information and Networks*, 2011, pp. 95–102.